STOCKHOLM SCHOOL OF ECONOMICS
Department of Economics
5350 Master's Thesis in Economics
Academic Year 2020-2021

Cyber-Speed and Scholastic Success:

The Effect of Internet Inequality on Education during the COVID-19 Pandemic

Adam Gill (41651)

May 29, 2021

Abstract

Technology is an integral part of modern society, but not everyone has equal access or quality. While this inequity in technology, known as the digital divide, affects many aspects of people's lives, education is one prominent area. Researchers have found that unequal access to and knowledge of computers, the internet, and other technologies can lead to differences in education outcomes. Within this paper, I further the research about the effect of the digital divide on education by looking at internet inequality in the United States. Specifically, I look at how the shift to online distance learning affects the education of American students with differing internet quality, as opposed to access. For my analysis, I exploit the natural phenomenon of the COVID-19 global pandemic, which forced in-person schools to shift to online learning in late March of 2020. Using two difference-in-differences specifications, I explore the effect of internet quality, proxied by county-level median download speed, on three education outcomes: ACT scores, AP test scores, and graduation rates. I find that there is a significant effect of internet inequality on graduation rates, suggesting that students with low-quality internet are disproportionately affected by the shift to online learning in 2020 relative to the previous years. I find no significant effect on ACT scores or AP test scores.

Keywords: Digital Divide, Internet Inequality, Education, Online Learning, COVID-19

JEL: I24, O33

Supervisor: Abhijeet Singh
Date submitted: May 29, 2021
Date examined: May 24, 2021
Discussant: Michelle Rudolph
Examiner: Anders Olofsgård

Acknowledgements

This paper would not have been possible without the help, advice, and support of many people. First, I would like to thank my thesis supervisor, Abhijeet Singh, for his guidance through the writing of this paper and his useful feedback at the various stages of the process. I am also grateful to Emil Bustos and Felix Schafmeister for their helpful comments about and advice on the direction and content of my thesis. Additional thanks go out to Daniel Evans and Elena Clemente for providing me with useful notes and criticism during the writing process. I would also like to offer great thanks to my parents, Douglas and Carol Gill, for their continued support and advice while writing my thesis and through the entire master's program. Finally, I would like to thank all the people who offered me support and help throughout the thesis process. While I cannot personally name everyone who supported me here, special thanks go out to Albert Planting-Gyllenbåga, Hannah Ritchey, Michelle Rudolph, and Simran Pachnanda, all of whom provided continuous support despite suffering the brunt of my complaining when things were going wrong.

Contents

List of Figures

List of Tables

1 Introduction

In the modern era, technology has become a dominant element of many people's everyday lives and it is often taken for granted. However, not everyone has equal access to technology and one of the central questions concerning inequality revolves around the unequal access to and knowledge of technology. This "digital divide", a common nomenclature for the inequality related to technology, arose with the innovative booms of the 1970s and grew into a problem that began plaguing policy makers and researchers in the 1990s as the diffusion of personal technology and the internet expanded rapidly. To this day, countries with seemingly high levels of technological use, like the United States, still suffer from unequal access to technology. Even as access to technology and the internet, defined as a dichotomy of "haves" and "have-nots", becomes more ubiquitous, the quality of that access still differs between different socioeconomic and geographic groups. This leads researchers to ask a central question: How does this inequity in technology affect an individual's academic, economic, and societal future?

Within this paper, I look at the United States (US) and the effect of one aspect of the digital divide - internet inequality. However, unlike past literature which focused on whether or not people had access to the internet, I look at inequity in the quality of internet access. While recent policy changes and an increasing dependence on the internet is reducing the number of households with no internet access, there still exists a wide divide in the quality of the internet provided. As can be seen in Figure 1, there is still a wide gap in internet quality across the country when looking at county-level median download speeds, with the inequality stretching across states and rural-urban regions. Counties also have vastly different median internet speeds with ranges from 2 Mbps to 1000 Mbps, both relatively far extremes from the average fixed broadband download speed in the United States of 182.69 Mbps (Speedtest, 2021). Additionally, even with internet access being more widely available, there are still a number of students that lack internet access at home. While 88% of students aged 3-18 had access to the internet at home via a computer, according to a 2018 government report, 6% of students only had access at home through the use of a smartphone and 6% of students had no access to the internet at home at all (National Center for Education Statistics, 2020).

This inequity in quality is likely to have widespread effects on many aspects of society, among them education. As past research has shown, differing levels of access to the internet, whether at home or at school, has impacted education outcomes such as test scores, enrollment, and dropout rates (Attewell and Battle, 1999; Aydin, 2021; Fairlie, 2005; Mbunge et al., 2020). These effects on education are likely to have been exacerbated with the onset of the COVID-19 pandemic in 2020 as schools shifted to online distance learning for at least a portion of their school year. Since distance learning generally requires attending virtual classes with the use of livestream video services such as Zoom, the quality of internet at home becomes even more important. Worse internet quality could result in missing important information and generally receiving a poorer education than peers with better internet connections. With this in mind, I focus the analysis of this paper on one central research question: *How does the shift to online distance learning affect the education of American students with differing internet quality at home?*

Figure 1: Map of internet quality across US counties

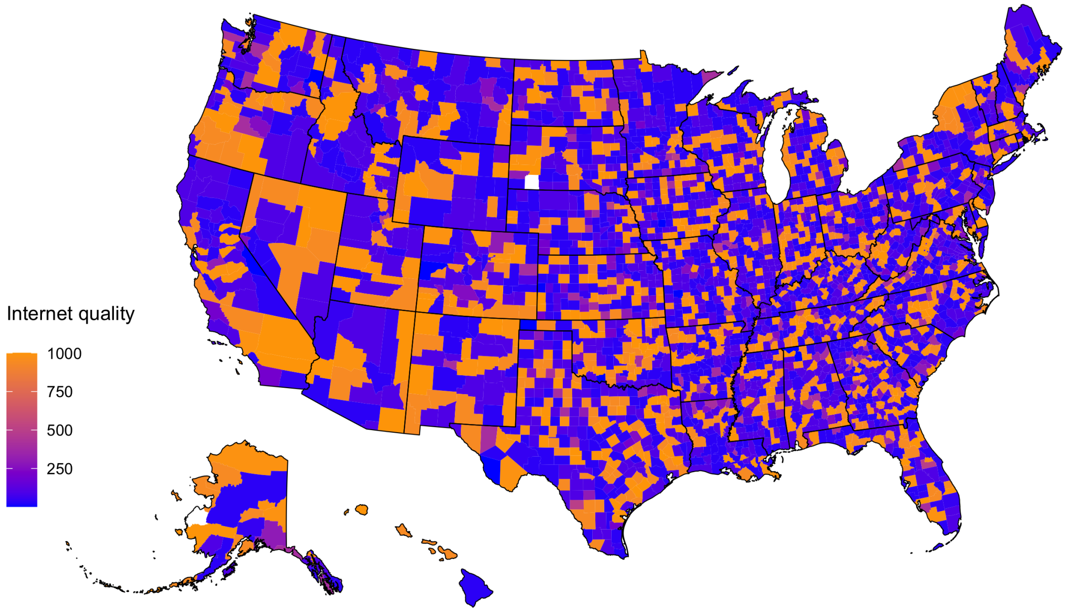

Notes: In this figure, blue represents areas with low-quality internet while orange represents areas with high-quality internet. "Internet quality" is proxied by county-level median download speed and is measured in megabits per second (Mbps).

Using a difference-in-differences (DID) analysis, I attempt to answer this question by looking at the effect that the inequity in the quality of internet access has on education outcomes when schools shifted to online distance learning in 2020. Specifically, I test the hypothesis that the education outcomes of students from counties with with low-quality internet are disproportionately affected by the shift to online learning caused by the COVID-19 pandemic in 2020 relative to previous years. I hypothesize that students with low-quality internet have a more negative or less positive effect on their education outcomes when compared to students with high-quality internet. To proxy for internet quality, I use median fixed broadband internet download speeds for counties in the United States, where lower speeds indicate lower quality internet. I run two DID specifications for my analysis. Specification 1 uses a generalized DID model in which the treatment variable is a continuous internet variable and uses county-level fixed effects to compensate for the time-invariant effects that were not partialled out. Specification 2 is a standard DID model using internet data binned into "low-quality internet" and "high-quality internet" bins. I run both of these specifications on each of my three education outcome variables: ACT scores, AP test scores, and graduation rates.

Through my analysis, I find some support in favor of my hypothesis as I get robustly significant results in the hypothesized direction for graduation rates when using the specification with the continuous internet variable. These results indicate that students from counties with lower quality internet are less likely to graduate in 2020 relative to previous years and compared to the trends of students from counties with high-quality internet. These lower graduation rates are likely caused by issues with internet connectivity and lower speeds. The poorer the internet quality, the more likely students struggle with attending classes and performing internet-based tasks. This may result in students receiving poorer grades, dropping out from school, or deferring graduation until after classes return to in person. While the results for graduation rates have significance, the results of my other education outcome variables are not significant.

My paper contributes to the literature in several ways. First, as far as I can tell, this is the first paper to use internet download speeds as a measure of internet quality when looking at education outcomes. This measure is particularly interesting as it relies on government data as opposed to the self-reported survey data that many other papers utilize. Additionally, due to the recency of the COVID-19 pandemic, there is very little research on this time in general and specifically in relation to education. My paper provides some early insights into the effect of the pandemic on education. Finally, this analysis adds additional evidence to the existing literature on the effect of the digital divide as well as providing additional insights into the growing literature of online learning.

My paper is structured in the following way: Section 2 consists of background information and previous literature about the digital divide; Section 3 consists of my research question and hypotheses; Section 4 includes a discussion of my data and a description of the education outcomes I have chosen; Section 5 outlines my econometric methods, specifications, and pre-trend tests; Section 6 includes the results for my various education outcome variables; Section 7 discusses the robustness checks on my results; Section 8 includes a discussion of the results; Section 9 discusses potential limitations with the analysis in this paper; and I conclude the paper in Section 10.

2 Background and Previous Literature

2.1 Defining the Digital Divide

Technological advancement has always been an integral aspect of human progress, but, beginning in the latter half of the 20th century, humanity entered into a new era of rapid technological change. Inventions such as the computer, the mobile phone, and the internet revolutionized all aspects of people's lives from economic opportunities to educational access to personal relationships. However, access to this technology has never been equitable as certain demographics, particularly minority, female, poor, and rural groups, disproportionately lack in both the access and education to take full advantage of it.

This technological inequality is often called the "digital divide", a term popularized amongst researchers and policy makers in the mid-1990s, despite the divide actually beginning in the 1970s with the rise of computers (Pierce, 2019). The digital divide is a global issue that acts as an umbrella term to discuss any technological inequality, whether it is between countries, socioeconomic groups, or even individuals. In the context of this paper, I focus on the micro aspects of the digital divide, looking at the effects between different groups within a country as opposed to technological inequality between countries.

While many technological innovations affected digital inequality, the emergence of the internet and its rapid adoption, in particular, seemed to widen the digital divide as the separation between "haves" and "have-nots" was not limited to computational power, but extended into information and knowledge inequality as well (Bucy, 2000). Academic, economic, and eventually social outcomes became more and more reliant on access to the internet. The emerging importance of the internet, as well as the wider availability of computers, brought about increased awareness and scrutiny about the digital divide among researchers and policy makers beginning in the mid-1990s (Pierce, 2019).

Early research into the digital divide focused on access to the various technologies. In 2000, Bucy looks at whether costs and expertise associated with the internet create an elite group of internet users. He finds that despite some increase in the diversity of internet users, it had not yet reached the point of being a widely used piece of technology. He continues by suggesting that this divide in internet users, especially along economic status since wealthier individuals were more likely to use the internet, could lead to a gap in knowledge and information. Another paper, by Strover (2001), looked into the geographic digital divide between urban and rural areas. She determined that there is inequality in computer access between urban and rural areas, with rural areas generally having much lower access. She also found that one reason for this inequality is due to the mismatch between rural residents and internet providers. Rural residents want access to the internet but only if costs are reasonable. However, due to the high cost of infrastructure development, it is only appealing to internet companies to expand into rural areas if they charge an inflated price. Looker and Thiessen, in their 2003 paper, build on Strover's paper by looking at this inequality not just between rural and urban areas, but also along gender and socioeconomic lines. Looking at students in Canada, they find that both rural and lower socioeconomic students are less likely to have computers at home and, for the lower socioeconomic students, this leads to the development of less positive attitudes towards technology. For rural students, access to computers in schools seems to compensate for the lack of access at home, at least in regards to attitudes towards technology.

While inequality still persist across countries, some authors have argued that in many developed countries, like the United States, the digital divide has been steadily decreasing. One paper, for example, looks at national surveys across the United States from 1995 to 2000 (Katz et al., 2001). Looking at the groups that had previously been underrepresented on the internet (such as women, minorities, and elderly), the authors find that there are swelling numbers of internet users among them and that the gaps between the groups are shrinking. While they are still below the proportional number of internet users expected, the growing trend would indicate a decrease in the digital divide. This becomes even more prominent when "awareness of the internet" is controlled for meaning that, of the subset of each group that knows what the internet is, the digital divide is even narrower.

While the digital divide may be narrowing in terms of having any kind of access to the internet, Attewell (2001) suggests that the digital divide is not disappearing, but has instead entered into a second phase. By surveying and synthesizing studies and surveys on the digital divide from the 1990s and early 2000s, Attewell determined that access was no longer the most prominent indicator of technological inequality. Instead, he suggested that, while surveys indicated that poor and urban students reported more computer use than suburban and affluent students, the way the computers were used differed greatly. According to Attewell, studies have found both that few children use home computers for educational purposes and almost all children use it for games or word processing. While these studies are outdated, they mirror more recent trends in computer use (Zhao et al., 2010).

Early experimental evidence for this second phase was provided by Hargittai (2002) who argues that the binary classification of having or not having access to technology is too constraining and overlooks some important aspects of the inequality. Through an experiment testing people's online skills, she found that, despite having access to the internet, people's ability to successfully complete tasks and the speeds at which they performed the tasks varied. This caused her to suggest that the digital divide be expanded beyond the dichotomy of "haves" and "have-nots" and to include the inequality in technological skill and learning. Gorski furthered the arguments posed by Hargittai. In his 2005 paper, Gorski advocates for a change in the digital divide discussion away from physical access and towards the inequity of use, particularly in regards to disenfranchised groups. Gorski finds that the use of the internet falls unfavorably along gender, race, and socioeconomic lines, which exacerbates the social and economic inequality these groups already face. While these disenfranchised groups have mostly overcome the divide in terms of have the physical capability of using the internet, they still generally lack the know-how to utilize this technology to its full potential.

In 2011, Wei and Hindman added additional support for Attewell's second-phase digital divide. These authors look at the relationship between the digital divide and the knowledge gap, specifically political knowledge. They found that the technology inequality between socioeconomic groups is not a result of unequal access, but is caused by inequity in the use of the internet, particularly in how information is gathered and communicated. As with previous research, they find that higher socioeconomic status (proxied by education level) is associated with more informational use of the internet. Additionally, they find that this difference in information attainment is more prominent in internet use than the use of more traditional media (like TV, radio, and newspapers). A more recent paper, by Subramony (2014), discusses how the digital divide is still a relevant issue despite the ubiquitous presence of technology in society. Like the previous papers, Subramony considers the digital divide a two-part problem with both access and technological ability as the key components. He wants to re-envision the digital divide as a social challenge rooted in socioeconomic institutions. Extending on previous work, the author suggests that the digital divide persists in part due to increased "cultural capital" of the traditionally entitled groups as well as the inequality in access, use, and ability.

DiMaggio and Hargittai (2001) also propose an extended version of the two-phase digital divide. Similarly to Attewell, they claim that the rapid diffusion of computers and internet access indicates that access, while still existent, had been becoming less of a determinant of the inequality. They suggest redefining "access to the internet" to being capable of performing necessary tasks and not just defining it as being connected to the internet. In order to do this, the term "digital divide" needs to be expanded. They reimagine the digital divide and include five additional aspects of technological inequality that need to be discussed and understood by researchers and policy makers: technical connections, autonomy of use, information and communications technology (ICT) knowledge and ability, social support, and purpose of use. According to the authors, all five of these characteristics are potential sources of technological inequity and could indicate a wider digital divide than previous authors considered.

2.2 Education and the Digital Divide

The digital divide has been and continues to be an important issue to understand and to find solutions for. Researchers have found that internet and technological inequality can negatively impact many areas of life for disenfranchised groups. One of the most affected areas is education and the digital divide's effect on education has been researched since the term was coined. One early paper studying the effect of technology access on education comes from Attewell and Battle (1999). They use data from 1988 to look at whether having a home computer affects reading and math scores for children. They find that having a home computer does have a positive effect on test results, but that this effect is disproportionately more beneficial to higher income than lower income students. The authors suggest that this is due to greater resources available to higher income families to ensure the computers are used for educational, as opposed to recreational, purposes. Huang and Russell (2006) also look at the amount of computer access students have and how this affects their academic achievement amongst students in Oklahoma. The authors find that there is a digital divide across socioeconomic groups that suggests the need for greater access to computers such as by improving the student/computer ratio. They also find that state test scores are affected by this unequal access, however, these results could be affected by other confounding factors. Aydin (2021) also follows up this research by looking at the effect of internet access on test scores in South Korea and Chile. In line with the results of previous research, he finds that greater access to the internet and increased computer literacy has positive affects on student test scores.

Others authors look at the connection between the digital divide and education, but from a different angle. Zhang et al. (2015), for example, also find that access to technology is positively correlated with better academic achievement, however they look at a different form of technology. They study access to mobile devices and the use of educational apps instead of access to the internet or computers like previous researchers have done. Fairlie (2005) uses access to computers at home, like previous research, but uses enrollment instead of test scores to measure the education effect. Fairlie looks at the impact of having computers at home and enrollment in high school (and, by extension, completion of high school) and he initially finds a positive relationship between owning a computer at home and being enrolled in high school to varying degrees. However, after including a number of controls, enrollment has a negative relationship with owning a computer at home among less educationally motivated families, which the author posits is the result of these families using computers more for recreational activities.

Other researchers have found that the perception of the importance of technology and a student's self-perceived ability to use it can affect academic achievement as well. Hargittai (2005) finds that self-reported levels of technological skills are generally overstated when compared to real measures of their skills. The author suggests that this overstatement is likely due to proficiency in the small selection of technological activities an individual usually utilizes despite a lack of skill in other technological areas. Peña-Lopez (2010) finds this gap in skill exists even among students entering college. He finds that there is a gap in technological skill with students entering college despite students feeling they are ready for college. He suggests that programs that increase access are not remedying the problem since they are misallocating resources away from solving the skill gap issue. One paper by Zhao et al. (2010), explores the DiMaggio and Hargittai framework in an academic setting and found that, even with access to the internet, differing socioeconomic and academic achieving groups tended to use computers in different ways and had different levels of internet confidence which creates another kind of digital divide. They perform some exploratory analysis looking at how internet self-efficacy and internet use affects test scores and they find that there is higher internet self-efficacy for students who use computers at home or at school is associated with better academic performance. This correlated with socioeconomic levels as students from higher socioeconomic groups generally spent more time on the internet and were more likely to be able to use the internet at home. They also find that students with lower levels of academic performance generally use the Internet differently that higher achieving students as lower academic achievement was correlated with using the internet for more leisure activities. However, since this is an exploratory analysis, no causal relationships could be determined. Buzzetto-Hollywood et al. (2018) find that this is still an issue, particularly for minority groups, in their more recent paper. They look at the "technological readiness" of students entering their freshman year of college at a traditionally minority-serving university and find that these students tended to be less proficient at academic-related computer skills such as creating presentations and making spreadsheets despite their prowess at using other computer applications such as internet communication and social platforms. Another paper by Heerwegh et al. (2016) may have a potential solution to the skill gap issue. They find that students who are exposed to computers earlier in life are both more likely to find computers useful and are more likely to have better computer skills.

The way teachers utilize technology is another component researchers look at in regard to the digital divide. Angrist and Lavy (2002) look at the introduction of computers in Israeli schools and they find that more computers increased the amount of technology-aided teaching, but this did not improve students scores. They suggest that differences in teachers' skills at incorporating technology into the classroom was likely an issue. Warschauer et al. (2004) furthered this idea by looking at how teachers use technology across socioeconomic groups. They find that teachers in lower socioeconomic schools are less proficient at incorporating technology into lessons and that this is likely exacerbating the inequity between groups.

2.3 Online vs. In-person Learning

Another growing stream of literature involves looking at the differences in student outcomes between traditional in-person classes and a growing number of online classes. Researchers have found various heterogeneous effects lead to differing results between in-person and online classes. One recent paper looking into live lectures compared to online lectures is by Figlio et al. (2013). In their paper, the authors use experimental evidence from university students in a microeconomics course who were randomly assigned to one of two lecture forms: live lectures or online lectures. The authors find that there is a slight cost to education outcomes when students learn through the online lectures compared to the in-person lectures, however, these results are not statistically significant. They do find that there is some heterogeneity in the results, however, with Hispanic students, male students, and lower achieving students having more positive results when attending live lectures than online lectures. Another paper by Bettinger et al. (2017), provides additional evidence about the difference between in-person and online classes. These authors use an instrumental variable approach to look at how students' grades are affected by taking online compared to in-person classes. Using a nation-wide college network that has both physical campuses and an online presence, the authors compare students' results across identical courses when they are taken online or in-person. The authors find that grades tend to be substantially lower for students taking courses online with an even larger negative effect on students who had lower GPAs ex-ante.

Additional evidence about online teaching also comes through looking at varying levels of hybrid teaching, a class format that combines elements of in-person and online teaching together. One paper that looks at the effect of hybrid teaching is by Joyce et al. (2015). These authors look at whether the amount of time spent physically in class affects a student's results. They randomly assign students to either a normal twice-per-week class or to a compressed once-per-week classes with both sections receiving the same online material that covers the entire scope of the class. The authors found that students in the traditional setting, where they had more classroom time, performed better on the midterm and final exams, however, only the difference in scores on the midterm exam is statistically significant. Also looking into the difference between traditional and hybrid classes are Haughton and Kelly (2015). In their paper, the authors look at whether students who view lectures online first and then have a face-to-face meeting with the professor (flipped hybrid model) perform differently than those students in the traditional in-person lecture style classes. The authors find that, for one of the semesters, students in the flipped hybrid classes performed better on the common final exam than students in the traditional class, but the results were insignificant for the other semester despite similar structure of the courses and similar demographics of the sample. The authors also find no significant difference in the other course-related outcomes such as semester grades or perceived course rating between the flipped hybrid and the traditional classes in either semester.

Cacault et al. (2021) provide further discussion on the effect of online vs. in-person learning. In their paper, they ran an experiment where each week some students were randomly given access to the online, live-streaming version of the lecture while still having the opportunity to attend in person. The students who were not provided with access to the platform that week only had the option of attending the class in person. Through this design, the authors are able to look at a hybrid structure compares to an in-person structure on a lesson by lesson basis as well as how having access to online affects attendance. They find that students generally choose in-person classes even when given the option online, and tend to only opt for online over in-person classes when there are high costs to attending in person. Additionally, they find that there is a heterogeneous effect on exam question scores from attending online classes or in person. Specifically, they find that online classes generally negatively impacts the scores of lower-ability students while it has a positive impact on the scores of higher-ability students.

However, not all researchers found a difference in performance between online and in-person classes. A recent paper by Merkus and Schafmeister (2021) exploits a hybrid system devloped by a Swedish university due to COVID-19 to look at whether attending in-person seminars or online seminars affects student performance in a bachelor-level economics course. They find no statistically significant difference in scores between students who attended relevant in-person seminars compared to students who watched the seminar recordings online.

My paper further contributes to this stream of literature by looking at another potential avenue of heterogeneity that could lead to differing results between in-person and online classes - internet quality. Differing education results that occur when shifting to online classes could be the result of poorer internet infrastructure and internet access, which would create a gap in the quality of education that would not be present during in-person classes. Since the structure of online classes and the programs and tools utilized to facilitate online learning are generally dependent on certain amounts of internet access and speeds, students in areas with poorer internet quality are likely to be more disadvantaged. This difference between online and in-person classes has become even more critical during the 2020 pandemic as schools were required to physically shut down and distance learning for classes occurred across most districts in the United States to varying degrees.

2.4 COVID-19 and the Digital Divide

COVID-19, the pandemic that spread throughout the world in the early months of 2020, caused much disruption and altered people's routines and daily lives. National and state governments issued various protocols such as lock downs, wearing masks, and social distancing in order to attempt to mitigate the effects of the virus. While many aspects of life were altered during the pandemic, education was among the most affected. Statewide closures of public high schools in the United States for at least part of the spring of 2020 were put in place in all states except for Wyoming and Montana and even those two states saw many schools closing at the district or county levels (EducationWeek, 2020). With school buildings closed, students still needed a way to learn so distance learning became a prominent tool to use. Most schools in the United States shifted to online distance learning where students, through conferencing apps like Zoom, attended virtual classes. This shift to online learning is expected to have an impact on students' education, especially early in the pandemic when the technology was new and teachers had little time to adjust their lessons plans.

Due to the recency of the pandemic, however, research on the effect of COVID-19 on education is relatively scarce. However, there have been research studies that look at how this disruption affected education and whether or not the digital divide played a role. Two papers, both by a similar group of authors, look at the whether the shift to online learning caused by the pandemic has an effect on university students in Zimbabwe (Mbunge et al., 2020; Mbunge et al., 2021). In one of the papers, they find that there seems to be an increase in the dropout rate of students due to COVID-19 and that this increase seems to be partially due to issues with internet connectivity (Mbunge et al., 2020).

My paper adds to the research on the digital divide by looking at how inequity in internet quality affects education outcomes during the pandemic and it contributes to the literature in a few key ways. First, I measure internet inequality using quantitative government data, which differs from the self-reported survey data that much of the previous research utilizes. While this prevents me from looking at household-level access to data, it allows me to look at a broader geographic area and to avoid any selection or response bias usually present in survey data. Additionally, I look at internet inequality in a different way. I avoid the often utilized dichotomy of "access" and opt for looking at internet inequality in terms of internet speeds. Based on my data, many counties are recorded as having median internet speeds lower than the optimal or even the minimum amount needed to run certain commonly used programs. This suggests there is an additional level of internet inequality beyond just access. To the best of my knowledge, this is the first paper to look at how inequality in internet speeds affects education outcomes. Finally, this paper offers an early contribution to the literature on the effects of COVID-19 and distance learning on education.

3 Hypotheses

In this paper, I look at how internet inequality affects education outcomes. However, unlike the stream of research that focuses on the digital divide in technological skills and knowledge (see Attewell, 2001; Hargittai, 2002; Wei and Hindman, 2011), my paper looks at internet inequality in access. Using aspects of DiMaggio and Hargittai's framework of internet inequality, I expand the concept of access beyond the "haves" and "have-nots" dichotomy of early research (see Pierce, 2019; Strover, 2001), and instead look at access in terms of the quality of internet provided. While there is still a not-insignificant portion of US households lacking in any internet connection at home (National Center for Education Statistics, 2020), the 2019 internet data from the United States Federal Communications Commission (FCC) I use in this paper shows that every county in the United States has at least one area with access to broadband internet. However, having some ability to connect to the internet does not mean that internet inequality has disappeared. DiMaggio and Hargittai (2001) suggest that *technical connections* is another aspect of the digital divide and this includes the hardware and means to access the internet. According to these authors, having different types and quality of internet access is a further extension of internet inequality in access. Following their argument, I utilize differences in internet speed as a measure of internet quality to explore the inequity in technical connections.

I look specifically at the effect on education. In the academic context, lower internet speeds can lead to issues of applications not working, poorer quality connections to online classes, and longer time spent downloading and uploading assignments. Prolonged internet issues for academic purposes could lead to poorer performance in aggregate education outcomes, such as standardized test scores or likelihood to drop out of school (Aydin, 2021; Fairlie, 2005; Mbunge et al., 2020). These factors are likely exacerbated by the increased reliance on the internet during the COVID-19 pandemic when most schools switched to online learning. This leads to the central research question of my paper:

Research Question: *How does the shift to online distance learning affect the education of American students with differing internet quality at home?*

In order to answer my research question, I take advantage of a natural phenomenon that took place in 2020 - the COVID-19 pandemic. In the early months of 2020, a worldwide pandemic reached the United States with cases beginning to sore in March and early April. This led to schools closing their physical locations and shifting to online distance learning (EducationWeek, 2020). This large scale shift online for schools that had always been in-person allows for a comparison of education outcome between the 2020 online version and previous in-person years. Using education metrics that remain fairly consistent between years and across the country allows for a comparative difference-in-differences analysis between groups with different internet quality across different years. My hypothesis narrows down the specific scope of my paper to the changes in education outcomes caused by shifting to online learning during the pandemic. Therefore, I have developed one core hypothesis to test the effect:

H1: *The education outcomes for students in counties with lower-quality internet are disproportionately more negatively or less positively affected than for students from counties with higher-quality internet during the 2019-2020 school year compared to previous school years.*

For the purposes of this paper, I proxy "internet quality" using a measure of internet download speed. Therefore, counties with lower-quality internet are counties with lower internet download speeds and counties with higher-quality internet are counties with higher internet download speeds. Additionally, my hypothesis does not make any claims on the overall direction of the effect of 2020 on the education outcomes. Instead, it is worded as "disproportionately more negatively or less positively affected" which suggests that, regardless of the overall direction of the effect for 2020, students from counties with lower-quality internet perform comparatively more poorly in 2020 over previous years than students from counties with high-quality internet. This would mean that, should the effect in 2020 be negative, students from counties with low-quality internet would have a greater magnitude effect in the negative direction than students from counties with high-quality internet. Should the effect in 2020 be positive, students from counties with low-quality internet would have a lower magnitude effect in the positive direction than students from counties with high-quality internet. Additionally, students from counties with low-quality internet could have a negative effect while students from counties with high-quality internet could have a positive effect. All three of these scenarios would indicate support for my hypothesis.

In my hypothesis, the argument that students with lower quality internet have a more negative or less positive effect from the shift to online learning comes from the literature. Numerous authors have shown that education outcomes are affected by worse internet at home (see Attewell and Battle, 1999; Huang and Russell, 2006; Zhang et al., 2015). Some authors even look specifically at the shift from in-person to online learning and find that there are heterogeneous effects on students when changing between in-person to online, specifically with groups like Hispanic students, male students, and low-ability students worse off (Cacault et al., 2021; Figlio et al., 2013). The general consensus in the literature seems to indicate that less access to technology or access to lower quality technology has a negative academic impact on those students compared to students with access to high-quality technology. Therefore, I structure my hypothesis with the argument that students with lower quality internet are worse off from the shift to online learning than their peers with access to high-quality internet.

While I would ideally like to look at the overall effect of internet inequality and the shift to online learning on education, there is no aggregate variable that measures a student's overall academic achievement. Additionally, curricula and academic standards differ across states, counties, and schools, making comparison between them difficult. Because of this, I have chosen to proxy education outcomes with three academic variables that are standardized at the national level allowing for cross-county comparison. These variables are ACT scores, AP scores, and graduation rates. With these three education outcome variables, I form three sub-hypotheses from my main hypothesis:

- **H1a:** *ACT scores for students in counties with lower-quality internet are disproportionately more negatively or less positively affected than for students from counties with higher-quality internet during the 2019-2020 school year compared to previous school years.*

- **H1b:** *AP test scores for students in counties with lower-quality internet are disproportionately more negatively or less positively affected than for students from counties with higher-quality internet during the 2019-2020 school year compared to previous school years.*

- **H1c:** *The graduation rates for schools in counties with lower-quality internet are disproportionately more negatively or less positively affected than for schools from counties with higher-quality internet during the 2019-2020 school year compared to previous school years.*

Like my main hypothesis, each of my sub-hypotheses are based on evidence from the literature. The ACT and AP test scores hypotheses (H1a and H1b) suggest a negative relationship between low-quality internet and test scores. This comes from arguments in the literature that shows that greater access to the internet generally leads to higher test scores among higher achieving students (Aydin, 2021; Huang and Russell, 2006; Zhao et al., 2010). The shift to classes online should intensify this effect as it requires the use of online education tools, such as Zoom, which have certain internet speed thresholds for minimum and optimal levels of functioning (Zoom Help Center, 2021). Research has also shown that test scores are generally affected by the shift to online learning with heterogeneous effects between students of varying academic levels (Cacault et al., 2021; Figlio et al., 2013). My graduation rates hypothesis (H1c) also suggests a negative relationship between low-quality internet and graduation rates. This also comes from arguments in the literature which provide evidence that lower access to internet increases dropout rates in colleges and that having access to technology at home (such as computers and the internet) increases the probability that a student is enrolled in school (Fairlie, 2005; Mbunge et al., 2020). Graduation rates may also be similarly affected by the shift to distance learning as students with lower internet quality may have issues connecting or following classes and then might be less likely to attend online classes and more likely to drop out, reducing the graduation rate. There is some evidence to suggest this as students tend to prefer to attend classes in-person when given the option between online and in-person lectures (Cacault et al., 2021).

4 Data

For this paper, there are two main types of data needed: data on internet quality and data on education outcomes. The internet data comes from FCC governmental data and the education data comes from individual state departments of education. The education data consists of three variables of interest - two national standardized test scores and graduation rates. The two standardized test scores are ACT scores and AP test scores.[1] This data includes the COVID-19 period (2019-2020 school year data) as well as two pre-COVID periods (2017-2018 and 2018-2019 school years). All of the data is aggregated to the county level. Within this paper, I assume each county is a self-contained entity where students live and attend school within the same county. This is a reasonable assumption because my data consists of public schools within the United States where students are allocated to schools based on their home address. Therefore, apart from a few isolated exceptions, American students attend high schools in their counties. Along with the data for the main variables of interest, I also include a discussion on the data sources for the control variable, COVID-19 infection rate. The details of the data sources, collection process, and manipulation of the data are detailed below with internet data discussed first (Section 4.1), the education data of ACT scores

[1] Another college-readiness standardized test, the SAT, is excluded from this analysis primarily because the tests normally scheduled to be taken at the end of the school year (April, May, and June) were cancelled in 2020 due to the COVID-19 pandemic, leading to no available data relevant to the question posed in this paper. However, the omission of the SAT data should not lead to any holes in the analysis as it is the standardized test that is the least related to material learned in class since it had long been a standardized test based more on innate aptitude than on knowledge and skills gained during high school. Even with the 2016 shift towards making the test questions more closely related to the Common Core curriculum, it still lacks the subject diversity of the ACT and AP tests and is the test that is the least directly classroom material (Lewin, 2014).

(Section 4.2), AP scores (Section 4.3), and graduation rates (Section 4.4) discussed next, and the control variable (Section 4.5) discussed last.

4.1 Internet Data

My paper utilizes internet data to look at internet quality of counties in the United States. In order to look at differences in internet quality, I use internet download speed as a proxy. For my internet data, I use publicly available data from the FCC based on Form 477. Form 477 is a government required document that broadband internet providers must file twice a year. In the form, they denote which geographic areas they provide some amount of internet access to as well as the max advertised download and upload internet speeds for those areas. This is registered at the census block level. Census blocks are classified as having access to the internet if at least one household in the census block receives or has the capability to receive internet access at speeds exceeding 200kbps. “Advertised” download and upload speeds for the purpose of form 477 are the internet speeds that a user could reasonably expect to receive from the internet provider and is not based on the “theoretical capacity” (FCC, 2019). This suggests that “advertised” internet speeds provided in the data should be fairly similar to the actual internet speeds households receive.

For my internet data, I use the December 2019 data from Form 477, which is the data from just before the onset of the COVID-19 pandemic. Because establishing additional internet infrastructure is a time consuming process, I do not expect that there are significant differences in internet access or speeds in the two years proceeding my data nor in the several subsequent months (the time period of my education data). I use the validated data set as the base for my analysis as opposed to the raw data. In the validated data set, staff at the FCC peruse the data from Form 477 and fix any mistakes in reporting as well as removing extraneous data (such as census blocks that only contain water). Because of this, I use the validated data which is more likely to reflect the true conditions of the population and does not include areas that are lacking in internet access because they are uninhabited, which, if included, could lead to bias in my results.

Since the data is provided at a census block level, I aggregate it up to the county level. I use the county level since it is the smallest geographic area that is comparable across all of my data sets. This is because some of my education data is collected at the county level and is not able to be dis-aggregated. Additionally, even the education data that I collect at a smaller geographic level usually exists at the school-district level, which does not always coincide with the census blocks thus making these smaller geographic levels incomparable. For my internet data, I aggregate up to the county level using the unweighted county median. Because of the relatively similar population counts across census block groups (generally between 600 and 3,000 individuals per block), I chose to use a uniform aggregation method (Census Bureau, 2019). I also chose to use the median as opposed to the average download speeds in order to compensate for the skewness of the distribution of internet download speeds towards higher internet levels.[2] Therefore, I construct my

[2]Counties generally have little variation in internet speeds across their census block groups so using medians over averages is expected to have little effect on the results. However, I use average internet download speeds as an alternative specification to

internet download speed variable as the median max advertised download speed for each county measured in $Mbps$ (megabit per second). General summary statistics of the internet download speeds by state can be found in Table A.1.

Within the data, the continued existence of the inequity of internet quality can be seen. Looking at the cumulative distribution function of the median download speeds for counties across the United States (Figure 2), there appear to be three ranges of download speed that the majority of counties fall into: 30-50 Mbps, 90-100 Mbps, 920-1000 Mbps. Because of this, I use two versions of my internet quality data. First, I use internet download speeds as a logged continuous variable to see what effect there is on education outcomes as internet quality increases. I log my continuous internet variable in order to adjust it for the long and large tail area (> 900 Mbps).

Figure 2: Cumulative distribution function for US counties' median internet download speed

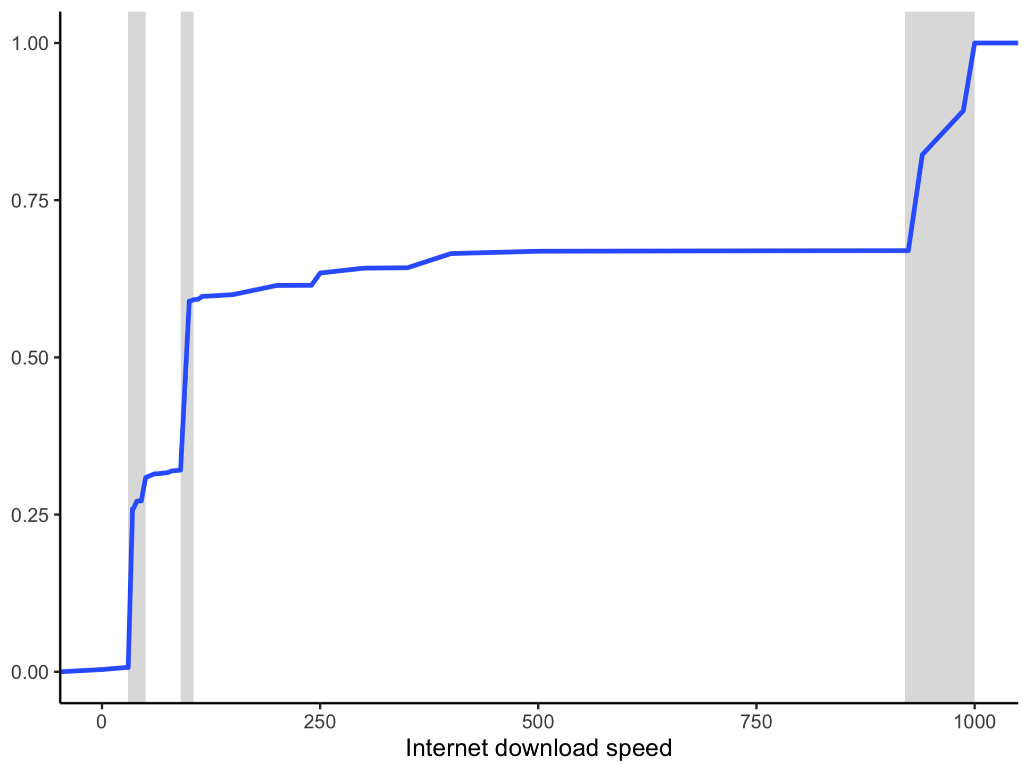

Notes: Internet download speed is measured in megabits per second (Mbps). The gray areas represent the three ranges of median internet download speed that are most common in US counties (30-50 Mbps, 90-100 Mbps, and 920-1000 Mbps).

Along with the continuous form of the variable, I also use a binned version of internet quality. While three distinct groupings of internet quality exist in the data, it does not make empirical sense to bin my internet data into these three groups. This is because, for the purposes of looking at internet quality and its

check the robustness of the results (see Tables B.1 and B.2 for ACT scores, Tables B.3 and B.4 for AP test scores, and Tables B.5 and B.6 for graduation rates).

effect on education outcomes, certain internet thresholds would perform similarly. Therefore, I would ideally want to find a cutoff speed that roughly divides the data between those with fast enough internet for online learning to function well and those with too slow of internet for online learning. Looking at official FCC guidelines, "advanced" (fast) internet for a household is defined as > 25 Mbps download speed and this speed is suggested for households where 2 to 4 people need to use high-usage programs, such as multi-person video calling, at the same time (FCC, 2020). However, 25 Mbps is likely still too slow of internet. One bipartisan group of senators is pushing for this definition of "advanced" internet to change from > 25 Mbps to > 100 Mbps (Bell, 2021). Yet, 100 Mbps is likely too high of a threshold for me to use as many students can probably succeed through online learning with lower internet speeds. Therefore, I choose a bin somewhere between these official and proposed official speeds.

Choosing the threshold speed, however, is not straightforward because there are many factors that can affect internet performance and the minimum speeds needed for a household including the number of people using the internet, the applications and activities being done, the number of devices connected, and the stability of the internet connection (Horaczek, 2020). Applications used for distance learning, such as Zoom, Google Hangout, and Skype require various amounts of internet download speeds to be functional. Zoom, for example, requires a minimum of 1 Mbps for group calling with recommended amounts of up to 4 Mbps (Zoom Help Center, 2021). While this may not seem to be a lot, these requirements refer to only the internet usage for Zoom. Often times, students also need to use the internet for other tasks concurrently like searching the internet, accessing school portals and websites, or checking emails, which increases internet speed requirements. An article on CNET provides basic and recommended usage requirements for various activities individually (Bolden, 2021). The basic usage requirements mentioned include email (1 Mbps), internet browsing (5 Mbps), social media (5 Mbps), and video calling (5 Mbps) with minimum recommended speeds being even higher for social media (10 Mbps) and video calling (10 Mbps). These are all activities that students could be using simultaneously or at least have open on their computer at the same time. This would suggest that, for one student, a minimum internet speed of 16-26 Mbps would be necessary.

This is further compounded, however, when there are multiple people at home all working or studying on the same network. "If you have several people working from home, including students, it can tax your connection" and voice and video chatting are especially taxing (Horaczek, 2020). This would likely be the case for most students working from home as it can be expected that they would have at least one sibling or parent who would also need to use video conferencing or other high-demanding programs for their school or work. Therefore, one additional person working or studying from home would need an equivalent amount of internet speed. For these reasons, it seems like internet speeds of 30-50 Mbps are probably conservative minimum requirements for household internet for students. Based on these considerations, I create a binned version of my internet quality variable with two bins referred to as the low-quality bin (≤ 50 Mbps) and the high-quality bin (> 50 Mbps). I use the binned version of this variable both when visually assessing the viability of the parallel trends assumption in my data (Section 5.3) and as one of my specifications for the difference-in-differences analysis.

4.2 ACT Data

Following the internet data, I obtain education data. The first of the education outcomes I look at is the results of the college readiness test called the ACT (American College Testing). The ACT is a multi-subject, nationally standardized comprehensive test that is meant to determine what a high school student has learned and how prepared they are for higher levels of education. One of the two prominent American college readiness tests (the other being the SAT), ACT scores are a measure accepted by American colleges and universities on applications. The ACT tests a wide range of subject areas based on the core curriculum of the United States so it is fairly reliant on classroom learning. The five subject areas are English (grammar and vocabulary), mathematics, critical reading, science, and writing. Students get individual scores in each of these subjects. English, mathematics, reading, and science are all graded on a 1 to 36 scale and the combination of these four scores gives a students composite score (also on a 1 to 36 scale). The writing section is optional and not factored into the composite score. Its score is the sum of two writing sections each scored on a 1 to 6 scale, giving an aggregate writing section score on a 2 to 12 scale.

For all of the education outcome variables, I was unable to get centralized data at a county or lower level. Therefore, I needed to collect the data individually for each state from each department of education or equivalent government body and aggregate them into one data set. However, not all states provide all the data I would like and not all states have the most recent year(s) available, so my analysis is based on the states which provide the adequate data for the desired years. Because of this, my analysis does not include all US counties and has a different set and number of counties for each of the education outcome variables. The limited number of states does restrict the external validity of my results as I do not have a broad enough sample of United States' states and counties for any of my education variables to make general claims about the entire country. However, I still believe the analysis in the paper contributes to the literature on the digital divide and education despite being limited in the generalizability.

For my ACT data, I have collected the school-district-level or county-level ACT results for seven states: California, Iowa, Louisiana, Minnesota, North Dakota, South Carolina, and Tennessee. Despite some loss of external validity due to the limited number of states, my sample of ACT scores data provides a nice cross-section of the United States' population. For example, it provides a geographically diverse sample with states from the West (California), Midwest (Iowa, Minnesota, and South Dakota), South (Louisiana), and Southeast (South Carolina and Tennessee). While there are geographic regions that are not included in my sample (most notably the Northeast), it still provides a varied geographic sample. It also provides a politically diverse sample with three states currently having governors from the Democratic party (California, Louisiana, and Minnesota) and four states having governors from the Republican party (Iowa, South Carolina, South Dakota, and Tennessee). While this does not mean my results have widespread generalizablilty, it does suggest that the results are not untowardly biased towards a certain region or political group.

For the states where the data was collected at the school-district level, they were aggregated to the county level and the ACT scores are the weighted county average results from the aggregated data. These aggregated averages are weighted by the number of students who participated in the ACT in that school district. The ACT data I use consists of various test score measures as well as the number of students taking the ACT. For each of these states, I have the results for three school years: 2017-2018, 2018-2019, and 2019-2020. This gives me the results of about 500 counties for each of the three years. For six of the states, I have data for the four main sections (English, math, reading, and science) individually as well as data on the composite scores. For my Louisiana data, however, I only have data on the composite scores, so this state is only included in composite score analyses. For the optional writing section, I only have data from one state (North Dakota), so I do not do any analysis on ACT writing scores. For six of the states, I obtained participation data directly from the state's department of education website. For North Dakota, however, they only provided a range of the proportion of students who took the ACT. In order to calculate participation numbers for North Dakota, I used enrollment data for the desired year and multiplied it by the conservative (lowest) proportion of the range.

Looking at the number of students taking the ACT, COVID-19 seems to have a negligible effect on participation. This is important, because, in order for the results of a difference-in-differences analysis to be meaningful, conditions that are not the variable of interest in the pre-treatment and post-treatment period should be as equal as possible. If the conditions are not equal, then the changes should, at least, have reasonable arguments that the treatment (in my case COVID-19) is not the cause of the changes. I check if there is any effect on the education outcomes that may be caused by COVID-19 in 2020 through a channel that is not the shift to online learning. In this situation, the number and composition of participation for my education outcome variables is one of these possible channels. However, if there seems to be no disproportionate change in 2020 compared to previous years then the DID analysis has a stronger argument.

Looking at the data presented in Figure 3, five of the seven states have relatively flat participation trends with similar number of students taking the test across all three school years. In California, there is a decrease in the number of test takers between 2019 and 2020, but that decrease is similar to the decrease seen between 2018 and 2019 indicating that it is probably a general time trend and not a COVID-19 related decline. These visual arguments are supported by the data as well. There is only minimal difference to the mean change in the number of students taking the ACT test for each state between the different school years for six of the states (Table 1). The only state that does not have a consistent trend across all three years is South Carolina, where participation in 2020 is similar to participation in 2019 (a mean decrease of only about 34 students), but where participation in 2019 saw a sharp decline from participation in 2018 (a mean decrease of about 383 students). However, it seems unlikely that COVID-19 would be responsible for this slowdown of the trend in the participation of ACT scores for South Carolina, therefore I keep it in my data.

Along with the relatively consistent participation numbers indicating COVID-19 had a negligible effect on the quantity of participation, it also seems unlikely that the composition of the participants would

Figure 3: Number of students taking the ACT in each year

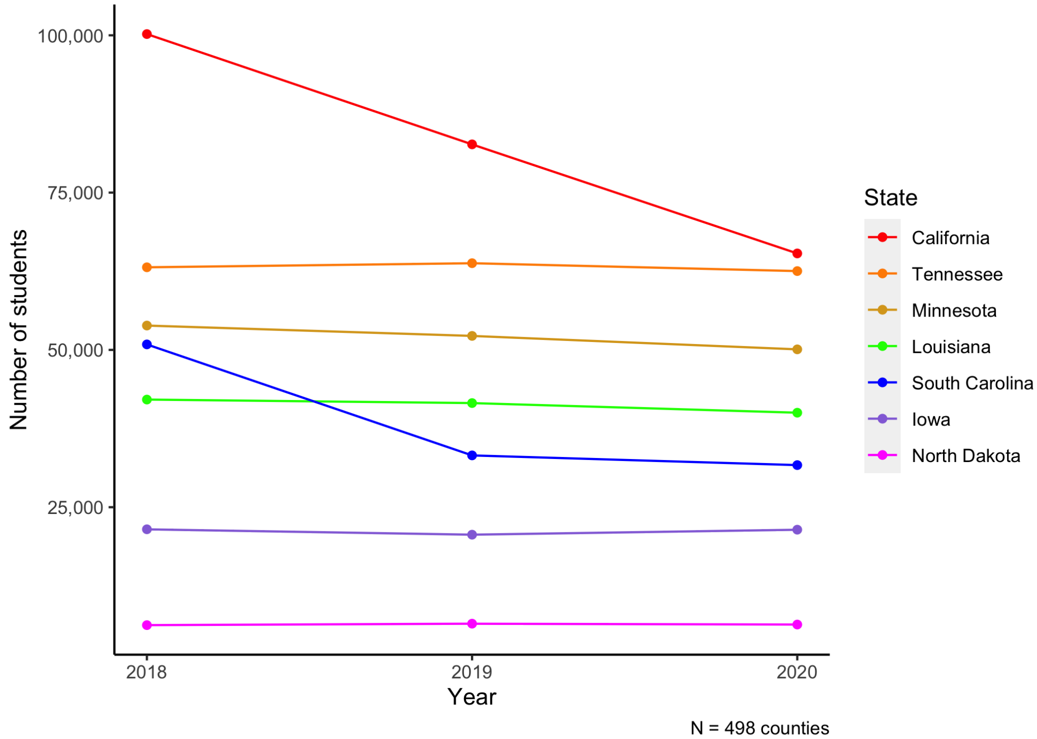

Table 1: Mean change in ACT participation

State	Change from 2019 to 2020	Change from 2018 to 2019	Difference in the changes
California	−298.90	−302.22	3.32
Iowa	8.03	−8.67	16.70
Louisiana	−23.94	−8.69	−15.25
Minnesota	−24.57	−18.92	−5.65
North Dakota	−2.96	4.73	−7.69
South Carolina	−33.57	−383.48	349.91
Tennessee	−13.20	6.87	−20.07

have changed either. While I do not have data on the "quality" of the students taking the ACT, there is no indication that selection bias would have taken place due to COVID-19. Looking at national ACT statistics, there appears to be little change in the students who take the ACT between 2018 and 2020 both in race and gender composition (ACT, Inc., 2018; ACT, Inc., 2020). While this is not county-level data, it seems reasonable that county-level trends in composition would mirror the unchanging national trends in composition over time. Additionally, taking the ACT is mandatory for most schools in two of the states within my data (Louisiana and North Dakota) and either the ACT or SAT is required for most schools in Tennessee. The use of ACT data in this analysis can be further motivated by empirical reasons. Because the ACT has subject based sections, it is likely to be more impacted by changes in how students learn from their teachers and the quality of time spent in class than similar college readiness tests like the SATs. Additionally, because it is a national standardized test, the results are comparable across counties and states.

However, there are a few potential drawbacks to using ACT data. While high school students in the United States can take either the SAT or the ACT, and both are accepted by almost all United States colleges and universities, each state tends to have a preferred test. Students generally take the test preferred in their state and these preferences tend to be regional so the sample of students who take the ACT may not be representative of the entire US high school population. The SAT and ACT are also designed to test different subjects, skills, and abilities so students may self-select into the test that matches their skill set. This could lead to different kinds of students taking each test. While these are both concerns to be considered and may reduce the external validity of the results, it does not diminish the results in the context of ACT test takers from the sampled states. Furthermore, while self-selection may be an issue when dealing with cross-sectional analysis of ACT scores, it should not have an effect on the difference-in-differences analysis used in this paper as long as the the self-selection does not fluctuate across years due to COVID-19 (which seems unlikely since participation and composition remained fairly consistent) and as long as any self-selection is not correlated with internet quality. Since college-readiness test selection would seems to be primarily driven by innate ability and geographic location, it seems unlikely that this self-selection would correlate with internet quality. However, I do not have the data to present formal evidence against any correlation, so it must be considered when interpreting the results.

Another potential concern with the ACT data is the disruption caused by the COVID-19 pandemic. Because of the rapid spreading of the virus in early spring across the United States, the initial test date on April 4th, 2020 was rescheduled to June 13th, 2020 (The Princeton Review, 2020). While this rescheduling may have caused some disruption to the testing schedule, I believe it has a negligible impact on my results. Despite the pandemic, the ACT test remained an in-person test at a testing site and retained the same structure and format as previous versions. Other than a few additional precautions due to the virus such as more space between students, fewer students per room, masks being worn when not seated at the desk, and greater availability of hand sanitizer, the process of the test was the same. Even the rescheduled test date is unlikely to have a large impact as the June 13th scores are included in the data for the same year and it does not seem to have affected participation rates (Figure 3). Additionally, because the ACT test remained in-person and no aspect of taking the test involves online activity, I expect that there is no correlation between internet quality and these changes to the ACT and therefore there is no effect on the error term. Because there seems to be no correlation and the effects of the changes to the ACT are uniform across the entire population, these changes are also unlikely to affect the analysis of this paper.

4.3 AP Test Data

Another education outcome variable I include is Advanced Placement (AP) test scores. The AP test is a standardized, subject-specific test taken at the end of certain high-level high school classes (referred to as AP classes). These classes are based on a nationally standardized curriculum and are designed to be at an analogous level to introductory college classes in that subject. At the conclusion of the class, the AP tests can be taken to show mastery of the course material and, if a high enough score is reached, some colleges and universities provide credits for AP classes in place of the equivalent introductory course. Different colleges and universities have different policies and minimum AP scores for accepting the tests for credit. The AP test is scored on a 1-5 scale with 5 being the best score. Each AP test is focused on one subject, but there is a wide range of potential AP tests from a variety of subject areas. Because AP tests are based on the material learned in a class, AP test scores are likely a good measure of how well students learned from their teachers and in their classes. Additionally, because it is a standardized test and the classes are based on a nationally determined curriculum, the scores allow for greater comparability across counties and states for schools that have the AP system.

For my AP scores, I have data from seven states: California, Kentucky, Louisiana, Massachusetts, North Carolina, South Carolina, and Wisconsin. Like the ACT data, I have a relatively diverse sample for my AP data as well. The AP data is geographically diverse as it includes states from the West (California), Midwest (Wisconsin), South (Louisiana), Southeast (Kentucky, North Carolina, and South Carolina), and Northeast (Massachusetts). While less politically diverse than the ACT scores data, the AP test scores data also provides some political diversity with five states currently having governors from the Democratic party (California, Kentucky, Louisiana, North Carolina, and Wisconsin) and two states having governors from the Republican party (Massachusetts and South Carolina). Similar to the ACT data, this is not meant to suggest that there is high external validity, just that the results are not too geographically or politically biased.

Similar to the ACT data, the AP test data comes from collecting data from individual state departments of education and, for all but one state, is at the county or school-district level. Data collected at the school-district level is aggregated up to the county level and the AP test scores are the weighted average scores for the county. For South Carolina, the data was only available at the school level. For this data, I first aggregate up to the school-district level using weighted averages, and then aggregate up to the county level in the same way as the other school-district-level data. For the AP data, all the aggregated averages are weighted on the number of AP tests taken.

While I would ideally like to construct my AP test scores variable using the absolute test scores by subject, the structure of the data I have access to has restricted the form that this variable can take in two main ways. First, while there are many types of AP tests that can be taken, I do not have access to test or subject level data for some states. For these, I only have some measure of the aggregated AP scores. Because of this, my AP scores variable is an aggregation of all of the tests taken, regardless of the subject. Despite the wide range of AP tests in the data, this aggregation should not be an issue since there is only a negligible change in the number of AP tests taken across the years and the distribution of the test subjects taken changed insignificantly (College Board, 2018;College Board, 2019; College Board, 2020a). Another data limitation I have is that, for several states, they present their AP scores as a percentage of students who who achieved a "qualifying" or "passing" score (a 3 or better on the exam) instead of showing the number of tests that achieved each score. While I would ideally prefer to run my analysis on the distribution of the actual scores achieved, this limitation in the data has led to my AP scores variable being constructed as a percentage of students with "qualifying" scores. For the states where the data is the full distribution of scores, I construct the AP scores variable by summing the number of tests that scored a 3, 4, or 5 and divide it by the total number of tests taken.

Apart from the structural issues in the data that restrict the format of my AP test scores variable, there are also some potential empirical issues with AP test scores that should be considered. First, there is no hard cap on the number of AP tests a student can take in a year. While each student can only take a specific AP test once, it is common for students to take multiple AP tests each year. While this multiple test taking could compound sample bias in my analysis should the make-up of the students differ across years, there is no evidence to suggest that the type of students differed.

Another potential issue in the AP test score data is the effect of COVID-19 on the tests. Similar to the ACTs, the 2020 pandemic disrupted the testing process for AP tests, however it did not change the date of the AP tests. They still took place on the scheduled dates, with all the tests being given sometime between May 11th and May 26th (College Board, 2020b). However, unlike previous years, the AP tests for each subject took place at the exact same time across the globe in order to reduce the chance of cheating. Along with the time change, the format of the AP tests changed. Instead of the normal two-hour exam with multiple choice and essay questions taken in schools the exam shifted to a 45 minute at-home, online exam with only essay questions (Hess, 2020; West and Johnson, 2020). Because it was taken at home, for 2020 only, the test required access to the internet to be taken which could lead to issues when trying to isolate the effect that internet quality from shifting to online learning had on the AP scores from any other effects due to these changes. However, even with the shift to taking the exam at home and online, I do not believe there should be a significant impact on the results as I expect any effects caused by this shift to be minimal in magnitude, relatively uniform across groups, or not correlated with the error term.

The first potential issue with the shift to the truncated, online AP exam revolves around access to the exam itself. Students in areas with low or no internet access may be unable to access the test or have trouble through the actual testing process. While this is definitely a concern, College Board (the organization that runs the AP tests) attempted to help remedy this situation by providing loaner devices and WiFi to schools and communities who requested it (Richards et al., 2020). This would help reduce the gap in the conditions of test taking, but would not impact the internet quality effect on online learning since it was temporary and installed shortly before the exams were to be taken. There also seems to be little to no change in the participation numbers of AP tests. Looking at Figure 4, there appears to be relatively consistent participation from 2018 to 2020 for six of the seven states I have data for, with California being the only state that sees a noticeable decline. However, even the decline for California is unlikely a problem as there is both no evidence that the decline is due to the onset of COVID-19 and the decline represents only a small decline in the percentage of the total tests taken (approximately 5%) which is even smaller when compared to the entire population of California high school students. Looking at the mean change in AP participation (Table 2), there is further evidence to support the minimal change in participation. Only two states have a difference in average county mean changes of more than 130 students (California and Massachusetts) and these mean changes are driven by large urban centers in both states. Any change in participation unrelated to issues with accessing the test is also less likely to be caused by COVID-19 as students choose AP classes to take at the end of the previous school year which would be prior to the onset of the COVID-19 pandemic. Along with little change in participation, the AP test saw little change in the composition of students taking the tests. While I do not have county-level data on the composition, the national test statistics show roughly similar student demographics for AP class composition and AP test participation composition (College Board, 2018; College Board, 2019; College Board, 2020a). Since AP participation and composition remained fairly consistent, the change in the format of the AP tests seems to have a negligible effect on who takes the AP test.

Figure 4: Number of AP tests taken each year

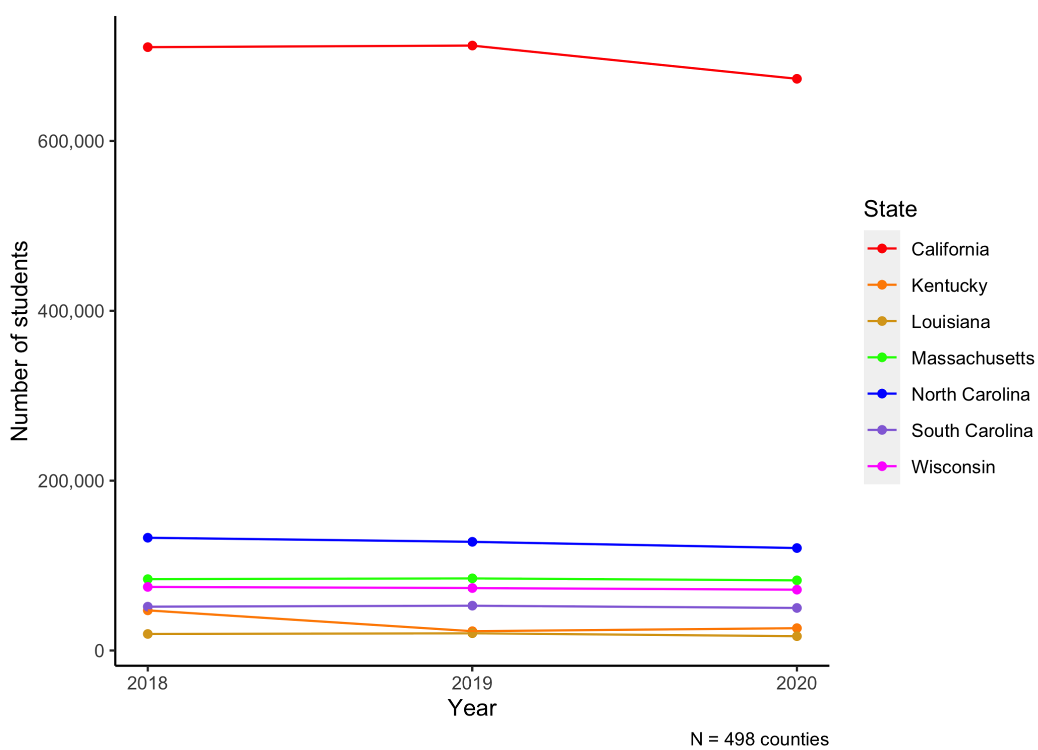

The timing of the exam is another disruption in the AP test process, however, this should not be problematic either. While each AP test was changed to be taken at the same time across all time zones,

Table 2: Mean change in AP participation

State	Change from 2019 to 2020	Change from 2018 to 2019	Difference in the changes
California	−699.5	34.2	−733.7
Kentucky	−62.7	−77.9	15.1
Louisiana	−98.9	30.4	−129.3
Massachusetts	−180.2	55.2	−235.5
North Carolina	−86.4	−55.5	−30.9
South Carolina	−82.0	35.4	−117.4
Wisconsin	−27.4	−19.0	−8.4

the timing of the exams were normal school times for students located within the United States, where the vast majority of the students in my dataset are located. Having the exams taken at the same time should only affect students from drastically different time zones of which international schools with AP programs would not be included in my data. The only international students that would be in my data would be ones enrolled in US high schools, generally as exchange students, and they would consist of a a very small portion of the student population. Therefore any effect from this is expected to be negligible. Additionally, any effect from this problem is independent from internet quality as it should affect international students uniformly and thus it would not correlate with the error term.

Other potential issues revolve around the test-taking process itself. There could be some concerns that having lower internet quality could affect the actual testing conditions. However, because of the design of the test, I do not think that this has any measurable effect on the results. The shift in format to one or two essays and the reduction in the timing of the test should not be an issue as it affects all groups uniformly and is uncorrelated with internet quality. Taking the test online, while directly connected to the internet, I expect to have negligible differences in effect across students with dissimilar internet quality. First, while the questions are online, internet speed should not be a large concern as each test only has one or two essay questions which are presented on one or two website pages. This means that there is minimal toggling between pages and limited amount of time spent waiting for pages to load. Therefore, even students with slower internet should not lose more than a few extra seconds compared to their peers with faster internet. Similarly, students were allowed to use the internet, as well as their notes, to search for answers during the test. If internet searching can lead to better test scores, then there would likely be correlation between the residuals of the analysis with internet quality as students with access to faster internet could search for information faster and thus have more time to write the answers. However, I do not believe this is an issue as College Board claims that the restricted time for the test and the nature of the questions limits the efficacy of searching on the internet (Kircher, 2020).

Additionally, students had the option to either type the answers into the online test or to write the answers and upload them as pictures. This further mitigates any internet issues around test taking as only limited computer skills are necessary for taking the test. However, due to the correlation between typing abilities and internet quality (see Aydin, 2021; Heerwegh et al., 2016), one potential problem is that students with better technological skills can type faster than they can write (see Mogey and Fluck, 2015). This would mean that these students could write more in the allotted time than they otherwise would have if handwriting and they could potentially write more than their peers who are handwriting. However, I do not feel that this is too problematic either, both because the short time period for the reduces the benefit of this faster writing and because the scoring of the writing is not based on length but based on the strength of the arguments. Therefore, short, well-written essays would score better than long essays with weaker arguments.

By far, the most likely problem to affect my results is issues with submitting the exam. Due to the short time between the onset of the pandemic and the scheduled AP test dates, College Board had little time to adapt the exams to an online format, and there were bugs in the system, particularly in regards to submitting the tests upon completion. There were approximately 10,000 students who failed because they were unable to submit their exams correctly or in a timely matter due to these issues which led to public backlash about the exams and a civil suit (Richards et al., 2020; Snouwaert, 2020). However, I do not believe that these submission issues should cause problems with my data either. While 10,000 students were affected by the submission issues, that was less than 1% of all the AP test takers in May 2020 (Richards et al., 2020). Since it is such a small percentage of test takers that were affected, it should have a minimal effect on my results. The effect on my results is further mitigated by the fact that College Board attempted to remedy the situation by allowing those who face submission issues to retake the exam for free in June and these results are also in my dataset (Snouwaert, 2020). While retaking the exam is not ideal and could lead to some bias in this small portion of the population due to differing testing conditions, I expect that it has a smaller effect than if all of these students had failed. Another reason for why this submission problem should have only a negligible impact on my results has to do with who was affected by submission issues. While internet quality did cause some of the submission issues with students with lower quality internet more likely to face submission issues due to lower internet speeds, another major submission problem was related to the format of uploaded pictures. Specifically, students using newer iOS devices (such as iPhones) to upload pictures of their exams found that the format of their phone pictures were not accepted by the AP test online portal (Chin, 2020). Students facing this type of submission issue were more likely to be from higher socioeconomic families (as it affected newer, more expensive technology), and higher socioeconomic families are more likely to live in areas with access to better internet quality (Wei and Hindman, 2011). Because of this both students with low-quality internet and students with high-quality internet were affected by the submission issue. While I do not have access to data on the exact percentage of students affected by each problem, the evidence suggests that students from counties of both high-quality and low-quality internet were affected to some extent, reducing any disproportionate effect on one of the groups meaning there would be an even smaller impact on my results from this submission issue.

While I do not believe these testing changes should affect the results of this paper to any significant degree, unequal testing conditions is a potential externality where the impact is difficult to discern. While there is some variation in testing conditions every year as students normally take AP tests in their schools and schools provide different testing environments, the COVID-19 pandemic could have created more disparity in testing conditions with students with low-quality internet more likely to face more difficult testing conditions. Since students had to take the exams at home in 2020, some students may face conditions that are unfriendly for test taking. For example, some students may have noisy households if many young siblings are studying from home and this could make it more difficult to take the exam (Richards et al., 2020; West and Johnson, 2020). While the shorter length of the AP test should minimize any effect from this, it should still be considered while interpreting the results of this paper.

4.4 Graduation Rates Data

The final education outcome variable I use is high school graduation rates. Graduation rates are measured as the percentage of students from a given cohort who meet the requirements to complete high school. I use four-year graduation rates which is the proportion of students who complete high school requirements in the standard four years. I include graduation rates because they provide a good balance to the standardized test data. Unlike standardized test scores which would generally have greater participation from and thus effect on the higher achieving students, changes in graduation rate would likely be caused by changes in the lower achieving students. Additionally, all high schools within the United States have graduation rates and graduation rates are calculated the same way across all schools making it good for comparison. The way graduation rates are calculated remains constant across years and is unaffected by the COVID-19 pandemic in terms of the construction of the variable.

As with the other education variables, the graduation rate data comes from collecting data from individual state departments of education at a county or school-district level. For states where I collect the data at a level smaller than the county level, I aggregate both the cohort size and the number of graduates to the county level and then recalculate the graduation rate by dividing the aggregated number of graduates by the aggregated cohort size. For the states that did not provide graduation rates but do provide number of graduates and cohort size, I calculate the graduation rates in the same way. All the graduation rates are the weighted average rates for the county and are converted to percentages. Along with graduation rate data, I have data on the size of the cohort for most of the states in my data. This data was also collected at either the county level or at the school-district level and then aggregated to the county level.

I have graduation rate data from 16 states:[3] Alaska, California, Colorado, Florida, Georgia, Idaho,

[3] I also have graduation data from Wyoming, however, I do not include these observations in my results because Wyoming is one of two states (along with Montana) that did not have statewide school closings.

Illinois, Kansas, Kentucky, Nebraska, New York, Ohio, Oregon, South Carolina, Washington, and Wisconsin. Similar to the other education outcome variables, I have a relatively diverse sample for my graduation rate data as well. This data is geographically diverse as it includes states from the West (California, Colorado, Idaho, Oregon, and Washington), Midwest (Illinois, Kansas, Nebraska, Ohio, and Wisconsin), Southeast (Florida, Georgia, Kentucky, and South Carolina), and Northeast (New York) as well as one state outside the contiguous United States (Alaska). The South is the only major geographic region that this sample does not have any states from. The graduation rate data is also a political diverse sample with nine states currently having governors from the Democratic party (California, Colorado, Illinois, Kansas, Kentucky, New York, Oregon, Washington, and Wisconsin) and seven states having governors from the Republican party (Alaska, Florida, Georgia, Idaho, Nebraska, Ohio, and South Carolina). The geographic and political diversity in my sample does not improve the generalizablilty of my results, but it does suggest that the results are not untowardly biased.

Similar to the ACT and AP data, I look at the change in cohort size over my three year time period to see if there is any potential COVID-19 effect that may result on changes in the sample for the 13 states I have cohort data for (with Florida, Illinois, and Kansas excluded because no cohort data was available for them). For standard cohort data, this would be trivial as cohort size would be determined when a class enters high school four years before graduation. This would mean that COVID-19 could not impact baseline cohort sizes as they would have been predetermined. However, the graduation rates presented by the states and the rates that I calculate utilize *adjusted cohort size* as opposed to baseline cohort size. Adjusted cohort size can change during the four years of high school as it alters the cohort size to compensate for non-academic reasons for students not graduating from a school such as removing students who transfer out or die and including students who transfer in during their high school years. While these changes are likely to be small, there still is a potential that COVID-19 could affect cohort size. Therefore, I check to ensure the trends don't change dramatically across years.

First, through visual inspection, there does not to be any significant changes in cohort size across the years for any state (Figure 5). This is expected as at least part of high school is mandatory in the United States so most high-school age children should be included in the cohort. This would also suggest that any small changes in cohort size is unlikely due to COVID-19 and is more likely due to shifts in the population of high-school age children. This is also supported by the quantitative evidence. In Table 3, the mean change in county cohort size by state is presented. Looking at the difference across years, most states see very little change, with the mean cohort for counties changing by less than thirty students for thirteen of the fourteen states presented. The only state to exceed thirty students is New York, which had a mean change in the county cohort size of 54.8 students. However, even this number is relatively small if considered in the context of the large populations some New York counties have (such as the five boroughs in New York City). Based on these arguments, it seems unlikely that COVID-19 has had any more than a negligible effect on adjusted cohort size.

Figure 5: Number of students in each adjusted graduation cohort

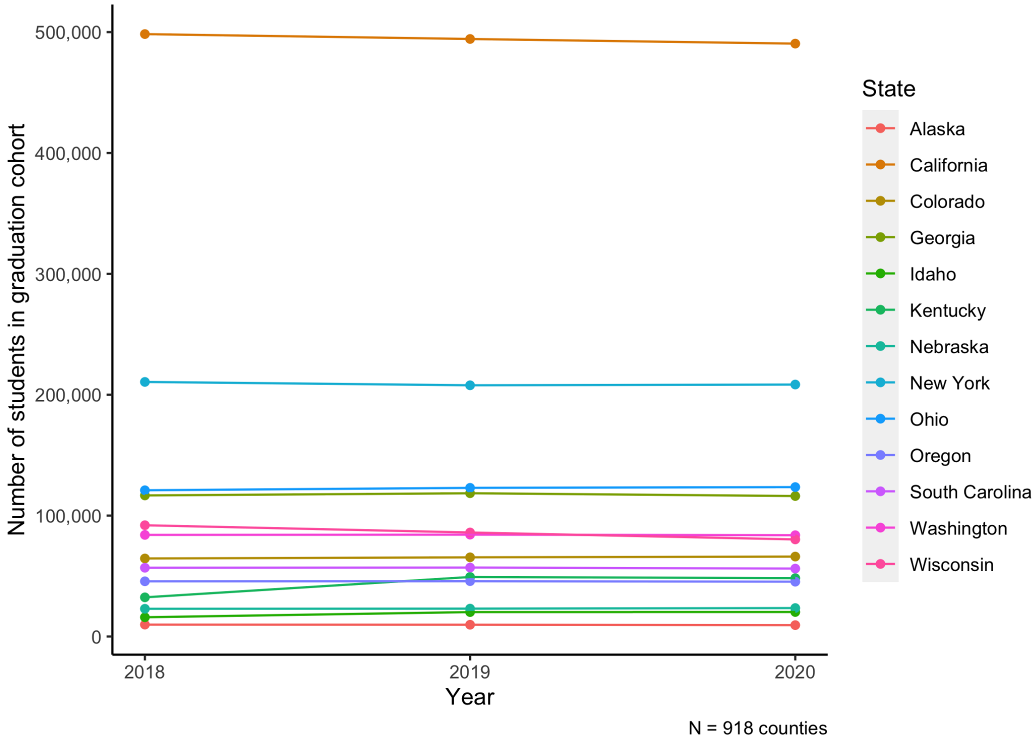

Despite the evidence against COVID-19 affecting cohort size, I acknowledge that there are some other potential limitations in my graduation rate variable. Unlike standardized test scores, which are administered and governed by a central body creating universal standards for all the tests taken, graduation rates are subject to more subjective bias and differences at the micro level. While federal and state education guidelines provide an overview of the expectations of a high school graduate, these are difficult to consistently enforce at the individual school level. Additionally, the requirements for passing classes are often left to the decisions of individual teachers which can lead to variation in graduation requirements. While any time consistent county-level variation would not affect my analysis (due to the difference-in-differences method), there is the potential for time inconsistent variation, particularly with the disruption caused by the COVID-19 pandemic. Some teachers may have lessened requirements or given extra leeway for students due to the trying nature of the lockdown. This would be problematic to my analysis if teachers offer disproportionate amounts of latitude to students based on their quality of internet. For example, if teachers systematically lessened requirements for passing classes for students who had poorer internet connection because they knew the lower quality internet made learning more difficult, this would bias my results. While I have not found any explicit evidence indicating a disproportionate shift in grading or graduation requirements, it cannot be ruled out and must thus be considered when interpreting the results in this paper.

Table 3: Mean change in graduation cohort size

State	Change from 2019 to 2020	Change from 2018 to 2019	Difference in the changes
Alaska	−11.5	−3.5	−8.0
California	−67.4	−71.1	3.7
Colorado	9.8	15.0	−5.2
Georgia	−14.8	11.8	−26.6
Idaho	4.3	5.6	−1.2
Kentucky	−7.6	1.3	−8.9
Nebraska	5.9	1.6	4.4
New York	10.1	−44.8	54.8
Ohio	7.2	28.3	−21.1
Oregon	−12.3	3.5	−15.8
South Carolina	−18.6	3.7	−22.4
Washington	−13.3	6.7	−20.1
Wisconsin	−79.6	−83.6	4.0

4.5 Control Data: COVID-19 Infection Rate

Along with my variables of interest, I collect data for my one control variable - county-level COVID-19 infection rates. This control variable comes from aggregated government data found on *USAFacts*.[4] I use two datasets from *USAFacts* to construct my COVID-19 infection rate variable. First, I use county-level COVID-19 data that includes a rolling sum of the total confirmed COVID-19 cases in the county beginning January 22, 2020. This data is originally sourced from the CDC (Center for Disease Control) and county and state health departments. The second dataset consists of county populations in 2020 which is originally sourced from the United State's Census Bureau.

These two data sets are combined to create my county-level COVID-19 infection rate variable. First, I find the average daily confirmed cases per county during a three-month period (April 1, 2020 to June 30, 2020). I chose this period because this overlapped the time that most schools shifted to online distance learning as well as overlapped the time when the ACT and AP tests were taken. With the average daily confirmed cases over the 91 day period, I then transform it into infection rate by dividing it by the county population from the other data set. I then multiplied it by 100,000 to get the infection rates per 100,000 residents, which is a common format used in medical science. This COVID-19 infection rate per 100,000

[4] *USAFacts* is a non-partisan, non-profit website that collects and aggregates data from various US government agencies for the purpose of increasing the accessibility of data and expanding analysis of government policies.

residents for each county is my COVID-19 control variable.

I include county-level COVID-19 infection rate as a measure for the intensity of COVID-19 and the effects of the pandemic on each county. I feel this control is necessary because areas were affected disproportionately by the pandemic with different numbers of people being infected. The varying intensity of the pandemic could have had non-uniform effects on the different counties in the United States leading to an impact on the education outcomes I look at that is not related to the shift to online distance learning. For example, students and teachers in counties with greater COVID-19 infection rates per 100,000 people are more likely to have to miss class, even online class, due to the virus, which is likely to impact education outcomes independent of the shift online. Additionally, students in counties with higher infection rates may face more emotional difficulties as it is more likely that they would personally know people who contracted the virus. A final reason is that schools were not closed for the exact same length of time. While school closures happened approximately at the same time (around late March), there was some variation (EducationWeek, 2020). Controlling for the infection rate, which would measure the intensity of the virus in an area, can also proxy for the length of time schools were closed as counties with higher infection rates were likely closed earlier and longer. For these reasons, I include the county-level COVID-19 infection as my control variable.

5 Method

5.1 Specification

For my analysis, I use a difference-in-differences method in order to measure the effect of inequality in internet speeds on education outcomes. DID analysis is a quasi-experimental design that attempts to create a pseudo-natural experiment using preexisting data. It does this by looking at a particular event and two groups - a group affected by the event (treatment group) and a similar group that was unaffected by the treatment (control group). It then compares the difference between the control group before and after the treatment and the difference between the treatment group before and after the treatment. In theory, this allows for a DID estimator that has group effects and time effects partialled out, allowing for the true effect of the treatment to be seen. However, DID is not a perfect estimator and caution should be taken when using it especially with regards to the underlying assumptions (Kahn-Lang and Lang, 2020), the pre-trends analysis (Roth, 2018), and potential endogeneity of the interventions (Bertrand et al., 2004). Despite these potential concerns, DID is a widely used econometric tool that, according to Angrist and Pischke (2010), is "probably the most widely applicable design-based estimator".

For the DID specifications in this paper, I use internet speeds as my "treatment" and the 2019-2020 school year as my "post-treatment period". Within my analysis, I have two pre-treatment periods (2017-2018 and 2018-2019 school years) and one post-treatment period (2019-2020 school year). While the differences in internet quality would have existed prior to the 2019-2020 school year, I argue that it would have had a much more negligible effect on education outcomes while schools were in person as was the case for the entirety of the 2017-2018 and 2018-2019 school years. However, once schools moved to online distance learning in the latter half of the 2019-2020 school year, differences in internet access would have had a much greater effect. I run separate regressions for each of my education outcome variables. Since there is no evidence of a pre-trend effect (see Section 5.3), I use a difference-in-differences specification that does not include a pre-treatment time trend. I use two different DID specifications for my analysis with a specification using a continuous internet variable denoted as Specification 1 and a specification using a binned internet variable denoted as Specification 2.

Specification 1

The first specification breaks from the traditional DID model structure as it treats internet quality, proxied by median download speed, as a logged, continuous variable. Because of this, it does not have a distinct divide between the control group and the treatment group, creating a difference-in-differences specification that is generalized. However, this DID specification with a continuous treatment variable is relatively common in the literature, with a number of papers using a similar specification, most notably the 2004 paper by Acemoglu et al. My specification is loosely based off of the model in the Acemoglu et al. paper, however, instead of including individual characteristics into the regression equation, I run my regression as a county-level fixed effects specification. This specification, using county-level fixed effects, is expressed as:

$$edu_{yc} = \beta_1 \cdot 2020year_{yc} + \beta_2 \cdot (Downloadspeed_c \times 2020year_{yc}) + a_c + \epsilon_{yc} \tag{1}$$

where y indexes the year and c indexes the county. The dependent variable, edu is the education variable of interest. the regressors are $2020year$, a dummy variable for 2020 (pandemic year), and the interaction of the 2020 dummy variable with the log of the median max advertised download speed. The county-level fixed effects are represented by a_c and the error term is represented by ϵ.

For this specification, the interaction term is an interaction between a continuous variable and a dummy time variable. This breaks from traditional DID specifications as there is no reference group ("control" group) built into the specification, so group effects are not partialled out by the interaction term. To remedy this, I use county-level fixed effects. The fixed effects model is a commonly used method for removing the unobserved, time-consistent effect of observations or clusters in panel data (Wooldridge, 2012). A fixed effects transformation is a method of running a regression where the effects of each unit in the fixed level are partialled out from the regression results prior to estimation. This is essentially the equivalent of running the regression with dummy variables for each unit in that level. For my regression, I use county-level fixed effects, which is the observation level of my data. In this case, running county-level fixed effects is equivalent to having a dummy variable for each county, thus partialling out all time-invariant effects. Because of this, both the non-interacted continuous internet variable and the control variable, COVID-19 infection rate, are co-linear with the county-level fixed effects, so they are dropped from this specification.

In the analysis of this paper, β_2 is the DID estimator, the coefficient on the interaction term between internet quality and 2020, the year of the pandemic when classes shifted online. In Specification 1 (Equation 1), this is the coefficient on the interaction between the 2019-2020 school year and the logged, continuous variable of median download speeds (*2020year*Download speed*). Unlike standard DID analysis, the β_2 coefficient in this specification has a different interpretation. Because it is an interaction between a continuous variable and a dummy variable, there is no standard "control" group used as a reference group. Instead, what this interaction term indicates is how having improved internet quality in 2020 changes differently from having improved internet quality before 2020. Essentially, β_2 on this interaction term represents the difference between the change in the education variable in 2020 compared to earlier years (2019 and 2018) for each one unit increase of logged internet quality. This means β_2 in this specification is not a fixed difference between groups and therefore this specification requires the use of a fixed effects model, unlike traditional DID analysis. Based on my hypothesis, I expect the sign of this coefficient to be positive indicating that, as the internet quality in a county increases, students are more likely to comparatively perform better on the education outcomes. This would be analogous to saying that counties with low quality internet are comparatively more likely to be affected negatively.

Specification 2

The second specification follows the traditional DID structure as it uses a dummy variable for internet quality where observations are binned into low-quality-internet and high-quality-internet groups determined by a internet download speed cutoff of 50 Mbps (discussed in Section 4.1). In this specification, the internet quality dummy variable equals 1 if the observation is in the low-quality group and 0 if the observation is in the high-quality group. This specification is expressed as:

$$edu_{yc} = \alpha_1 + \alpha_2 \cdot 2020year_{yc} + \beta_1 \cdot Lowquality_c + \beta_2 \cdot (Lowquality_c \times 2020year_{yc}) + \gamma_1 \cdot COVIDrate_{yc} + \epsilon_{yc} \quad (2)$$

where y indexes the year and c indexes the county. The dependent variable, edu is the education variable of interest and the regressors are $2020year$, a dummy variable for 2020 (pandemic year); $Lowquality$, a dummy variable that equals 1 if the median max advertised download speed of a county $\leq$ 50 and 0 otherwise; and the interaction of those two regressors. $COVIDrate$ is the control variable for county-level COVID-19 infection rate, γ_1 is the coefficient for the control variable, and ϵ is the error term.

In specification 2 (Equation 2), β_2 again represents the DID estimator and is the coefficient of interest. In this specification, it is the coefficient on the interaction between the 2019-2020 school year and the dummy variable for low quality internet (*2020year*Low quality*). In the second specification, this coefficient represents the fixed difference between the groups and the years with time invariant effects and group fixed effects partialled out. This is the DID estimator for the standard DID model, where the coefficient looks at the vertical difference between the actual results of the treatment group and the counterfactual results determined by the state of the control group. In this specification, the β_2 coefficient is the standard DID estimator since internet quality is binned into low-quality and high-quality groups. Therefore, β_2 shows if the expected mean change from pre-COVID-19 to post-COVID-19 is different between low-quality internet group (treatment) and the high-quality internet group (control) and the magnitude of the difference of the differences after the control variables are partialled out. Based on my hypothesis, I expect the sign on this coefficient to be negative. This would indicate that students in counties in the low-quality internet group are comparatively more likely to be affected negatively because of the shift to distance learning.

For the second specification, I use the control variable, COVID-19 infection rate. Because of the structure of difference-in-differences analysis, effects of time-invariant variables are partialled out of the coefficients of interest and thus do not need to be controlled for. Therefore, only variables that may change overtime differently for counties of varying internet speeds need to be controlled for. Because of this, I focus on one control variable in my analysis - COVID-19 infection rate. Since COVID-19 was a disruptive event that affected states and counties differently, I include a control variable that indicates the extent to which a county was affected by COVID-19. I use the infection rate as a measure for COVID-19 intensity. This control is included in Specification 2 regressions for all education outcome variables.

While I believe the control for COVID-19 infection rate is essential, I also intentionally limit the number of controls used. Standard DID specifications include implicit controls for time and group effects and the use of an interaction term reduces researcher degrees of freedom. Thus, as Kahn-Lang and Lang discuss in their 2020 paper, "with limited data, we face the usual risks of overfitting and loss of power when we try to address potential correlation with the error term by including further controls." Because of this, minimizing the number of controls included while still including all the necessary variables helps to prevent any overfitting problem.

5.2 Standard Errors

Within difference-in-differences specifications, there is no general consensus on the correct approach to take for formulating the standard errors and taking the wrong approach has the potential to generate incorrect standard errors that lead to misleading interpretation of the results (Wing et al., 2018). For example, in their 2004 paper, Bertrand et al. argue that the commonly used OLS standard errors, even corrected for within group-time correlation, is not adequate for DID analysis except under certain restrictive conditions as these specifications often suffer from serial correlation issues. Having a small number of clusters complicates the standard error analysis further. For example, Bertrand et al. propose a few alternative options, including block bootstrapping standard errors, however, this method requires a large number of clusters to be a viable option. Donald and Lang (2007) further this discussion by arguing that many of these standard-errors methods are even more inaccurate when the number of DID clusters is small. They claim that, when the number of groups is small, which is common in many DID specifications, standard errors formulated from more traditional methods are not normally distributed and thus alternative methods for obtaining the standard errors is needed. Cameron et al. (2008) also discusses issues with analyzing DID results with small number of clusters. Using Monte Carlo simulations, they show that the often-used cluster-robust standard errors result in over-rejection of the null hypothesis when the number of clusters is small. They also argue that the normal method for bootstrapping standard errors is also insufficient for correcting formulating standard errors when the number of clusters is small. Because of the limited data availability for the education outcome variables, there are a small number of state level clusters within my data (seven states for both ACT and AP data and 17 states for graduation rate data). Therefore, the more traditional methods of formulating standard errors, such as robust, cluster-robust, and pairs bootstrapping standard errors, are not adequate for my analysis. Various alternative approaches have been proposed for DID analysis with a limited number of clusters. One alternative is to use wild cluster bootstrapping. Cameron et al. (2008) find that wild cluster bootstrapping works particularly well for DID analysis with a small number of clusters, especially when compared to alternative approaches. The authors find that this method, a cluster generalization of the standard wild bootstrapping method, performs well even with as few as six clusters and that there is little loss in power after accounting for size.

However, despite a focus by some of these authors to utilize clustering of standard errors in DID analysis and Cameron et al.'s proposal to use wild cluster bootstrapping, the use of clustering has received mixed reviews more recently. While clustering standard errors had been common practice in samples where homogeneity is expected within clusters but heterogeneity is expected across clusters, arguments against their use have gained support in recent years. One pivotal paper looking at rethinking the clustering of standard errors is by Abadie et al. (2017). These authors argue that clustering in economics has become too common and has led to overly conservative standard errors. They suggest that clustering should be used only if one of two criteria is met - there is clustering in the sampling process or there is clustering in the assignment mechanism. Clustering in the sampling process is when a cross section of the population is created by randomly selecting from subsets of that population. Since the data in this paper consists of all counties where the data was available, there is no clustering in the sampling process. Clustering in the assignment mechanism occurs when the treatment effect is applied to grouped subsets of the data rather than applied to the individual observations. Within this paper, there is no need to cluster due to the assignment mechanism, since internet quality is usually assigned at the household level and even more general internet speed assignment (due to infrastructure) is assigned at the municipality level (Riddlesden and Singleton, 2014). Since both household and municipal internet assignment are lower geographic areas than the county-level data used in this paper, I do not need to cluster because of the assignment mechanism. Since the data in this paper does not meet either of the criteria posed by Abadie et al., I argue that there is no need to cluster the standard errors.

Following from this research, particularly the papers by Cameron et al. and Abadie et al., I use the wild bootstrapping method without clustering for the standard errors in my analysis. Unlike traditional statistical testing which compares test statistics to known distributions that are only approximately comparable to the true distribution, "bootstrapping" is a statistical tool in which simulated datasets are created by generating bootstrap samples (simulated draws) that mimic the distribution of the original sample and a test-statistic is then computed for each of these simulated datasets (Roodman et al., 2019). Bootstrapping can work well in many situations and the only main requirement for it to work is that the distribution of the test statistics must by asymptotic. The actual distribution does not need to be known or approximated ex-ante (MacKinnon, 2006). If the form of the heteroskedasticity of the standard errors is unknown, then wild bootstrapping, first proposed by Wu (1986), is one of the best methods available for adjusting the standard errors (MacKinnon, 2009). For the purposes of this paper, I specify the wild bootstrap method using the Rademacher distribution. This is one of the simplest specifications as the error term can only take one of two values (-1 or 1) with equal probability. Despite its simplicity, Davidson and Flachair (2008), using a series of simulation experiments, find that using the wild bootstrap with the Rademacher distribution is generally better and is never worse than using wild bootstrap analysis based on other distributions. They propose that using this form of wild bootstrap analysis should provide satisfactory inference in most practical contexts.

5.3 Parallel Trends Assumption and Pre-trend Analysis

Before running the difference-in-differences analysis, I look at the viability of the parallel trends assumption that is central to DID analysis. The parallel trends assumptions is the idea that, without the implementation of the treatment, the various groups in the analysis would have developed in the same way. In practice, this means that the difference between the various groups remains constant over time and that in the counterfactual scenario where the treatment never occurred, the groups would have continued to have a constant difference. This is called the "parallel trends" assumption because this constant difference between the groups is represented as parallel lines in a figure that maps effect over time. In the context of this paper, the parallel trends assumption assumes that, across the different levels of internet quality, the education outcome variables have consistent trends in years prior to the COVID-19 pandemic and would have continued on this same trend had COVID-19 never happened. If the parallel trends assumption holds perfectly, it would indicate that the difference in the effect (change in the slope after the treatment) is due entirely to the treatment implementation. Should the parallel trends assumption not hold, then it would be difficult to discern what part of the effect is due to the treatment implementation and what part of the effect was caused by some other exogenous effect.

While the parallel trends assumption is an important aspect of DID analysis, it can never truly be known if the parallel trends assumption holds. This is because it is reliant on the premise the the counterfactual situation (the scenario where the treatment never occurred) would have continued the same trends. Since, by definition, the counterfactual situation never occurs, it is impossible to say with complete certainty whether or not the effect trends would have remained consistent in the absence of the treatment. Despite a lack of certainty around the parallel trends assumption, there are analytical practices that can be utilized to test the viability of the parallel trends assumption. I use two main methods in order to explore the viability of the parallel trends assumption within my analysis by looking a the pre-treatment trends. First, I present a graphical representation of my data and visibly inspect the data for similar pre-trend slopes. Following the visible inspection, I use a Granger-type causality test to see if there are any treatment effects on my sample in the years prior to the implementation of the treatment.

Visual Inspection

The first method I use to assess the viability of the parallel trends assumption is visual inspection. Since my internet quality variable is a continuous variable, I bundle counties by median download speed in order to see if there are parallel trends. For these figures, I bin the internet data into the two bundles discussed in Section 4.1 with one bundle being the low-quality internet bundle (≤ 50 Mbps) and the other being the high-quality internet bundle (> 50 Mbps). Using the binned internet variable, I create event study figures for each of my education outcome variables to observe pre- and post-COVID-19 trends.

For the ACT data, I look at the pre-trends for the scores on each of the four main sections as well as two visualizations for the composite scores (one with and one without Louisiana data). The visual representation of the data is seen in Figure 6. The visualization of the composite scores (an average of the other four scores) overall has a positive change from 2019 to 2020 due to the inclusion of Louisiana which only provides data on composite scores. The visual representation of this data without Louisiana can be seen in the adjacent panel and has a negative change from 2019 to 2020, which matches the trend of the four individual scores. Looking at the pre-COVID trends in the ACT scores from 2018-2019 there seems to be similar time trends between the high- and low-quality internet bundles for four of the five scores - English scores, math scores, science scores, and composite scores. The similarity in the pre-trends for these variables provides evidence that the parallel trends assumption has some viability, which allows for greater confidence in the results of the DID analysis. Only the reading scores seem to have divergent pre-trends.

Looking at Figure 7, the time trends for the AP scores of the two groups seem to be less similar. From 2018 to 2019 the low-quality internet group is increasing and the high-quality internet group is decreasing. In fact, both groups have the same average percentage of qualifying AP scores in 2019 despite the low-quality internet group having a lower percentage in 2018 and the high-internet quality group having a higher percentage. Despite the lack of similarity in the pre-trends in this visual inspection, this does not necessarily mean that the difference-in-differences analysis for AP scores is invalid.

Looking at the event study figure for graduation rates, however, (Figure 8), both bundles seem to have similar pre-COVID trends. Both the low-quality and high-quality internet groups have positive time trends between 2018 and 2019. while the high-quality group might have a slightly greater increase, the pre-trends still seem similar enough to not obviously violate the parallel trends assumption.

Based on the visual inspection of the data for my three education outcome variables, there does not appear to be any obvious violation of the parallel trends assumption and even the variables with less similar pre-trends (namely AP scores and ACT reading scores) have support in favor of the parallel trends assumption from the Granger-type causality test (see below).

Figure 6: Average ACT scores

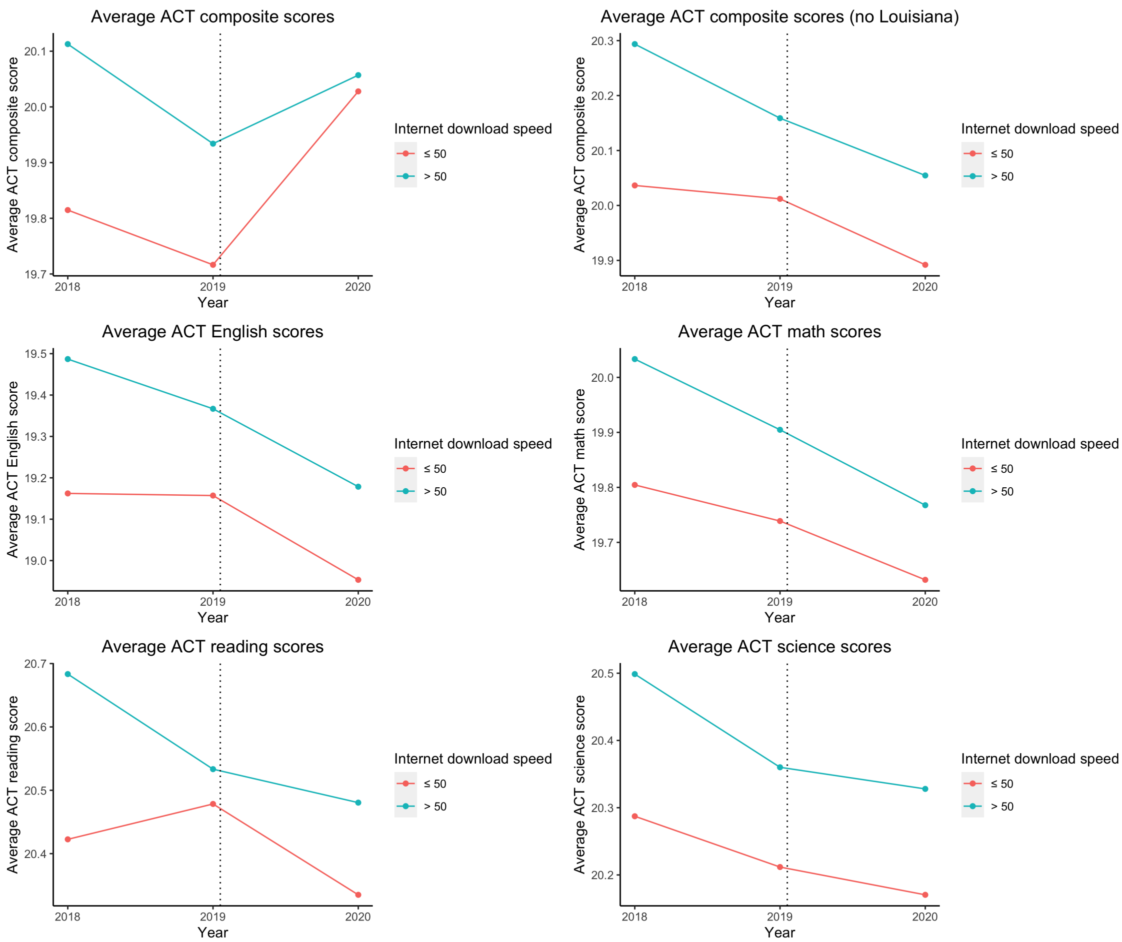

Figure 7: Average percentage of qualifying AP scores

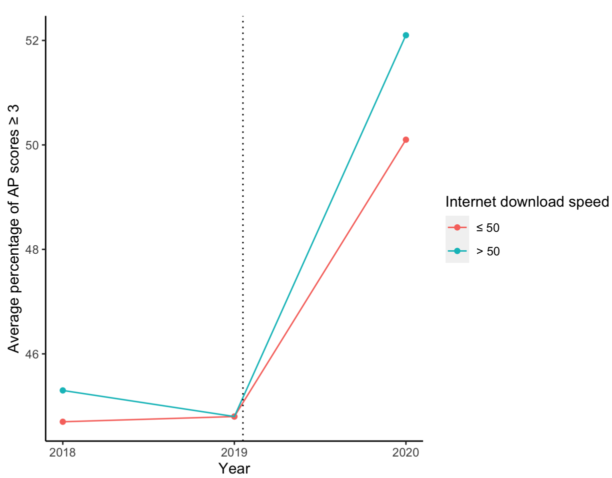

Figure 8: Average graduation rates

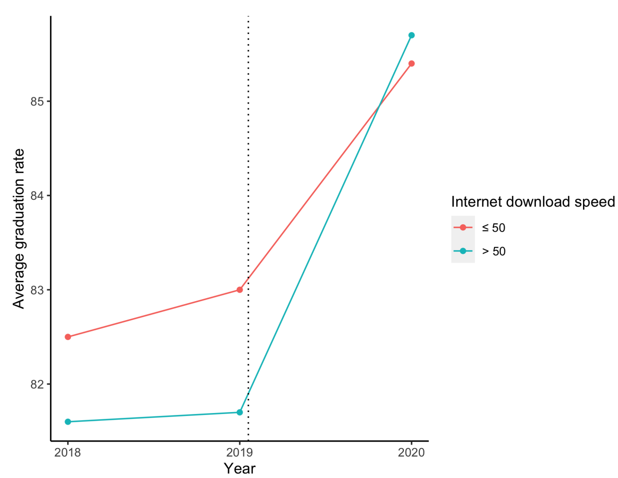

Granger-type Causality Test

Along with the visual inspection, I also perform Granger-type causality tests to explore the efficacy of the parallel trends assumption within my data. Granger-type causality tests look at whether or not there are treatment effects on the outcome variables prior to the assumed introduction of the treatment. In the case of this paper, I test to make sure that there is no treatment effect on test scores in the years prior to the pandemic. Specifically, I look to make sure that difference in internet quality does not have a significant impact on ACT scores, AP scores, and graduation rates between 2018 and 2019, the year before the implementation of my COVID-19 treatment and the shift to online learning. Should there be a significant effect, it would suggest that my base assumption that internet quality has a negligible effect on education outcomes before the pandemic does not hold which would invalidate the comparative analysis used in difference-in-differences specifications. Additionally, significance in the coefficient on the pre-treatment periods in the Granger-type causality test would indicate that the pre-trends are not similar and thus that the parallel trends assumption likely does not hold.

For the Granger-type causality tests, I use two specifications, one for a continuous internet quality variable and one for the binned internet quality variable. The first specification is for the continuous internet quality variable and I use the following specification:

$$\begin{gathered} edu_{yc} = \alpha_1 + \delta_1 \cdot 2019year_{yc} + \delta_2 \cdot 2020year_{yc} + \phi_1 \cdot Downloadspeed_c + \\ \beta_1 \cdot (Downloadspeed_c \times 2019year_{yc}) + \beta_2 \cdot (Downloadspeed_c \times 2020year_{yc}) + \epsilon_{yc} \end{gathered} \tag{3}$$

where y indexes the year and c indexes the county. For the variables, edu is the education variable of interest, 2019*year* and 2020*year* represent a dummy variable indicating 2019 and 2020 respectively, *Downloadspeed* is the continuous variable of median max advertised download speed, and ϵ is the error term.

The second specification is for the binned internet quality variable. In this specification, internet quality is a dummy variable that equals 1 if the observation is in the low-quality group and 0 if the observation is in the high-quality group. For these Granger-type causality tests, I use the a similar specification:

$$\begin{gathered} edu_{yc} = \alpha_1 + \delta_1 \cdot 2019year_{yc} + \delta_2 \cdot 2020year_{yc} + \phi_1 \cdot Lowquality_c + \\ \beta_1 \cdot (Lowquality_c \times 2019year_{yc}) + \beta_2 \cdot (Lowquality_c \times 2020year_{yc}) + \epsilon_{yc} \end{gathered} \tag{4}$$

where y indexes the year and c indexes the county. For the variables, edu is the education variable of interest, 2019*year* and 2020*year* represent a dummy variable indicating 2019 and 2020 respectively, *Lowquality* is a dummy variable that equals 1 if the median max advertised download speed of a county $\leq$ 50 and 0 otherwise, and ϵ is the error term.

For the purposes of these Granger-type causality tests, the coefficient of interest is β_1, the coefficient on the interaction between internet quality and 2019 (*2019year*Download speed*) in the first specification (Equation 3) and the coefficient on the interaction between having low-quality internet and 2019 (*2019year*Low quality*) in the second specification (Equation 4). For both specifications, in order to show that there are no pre-

Table 4: Granger-type causality tests for specification 1 (continuous)

	ACT scores						
	Composite	English	Math	Reading	Science	AP scores	Graduation rate
2019year	−0.164	0.071	−0.141	−0.037	−0.169	−0.163	0.482
	(0.484)	(0.631)	(0.497)	(0.526)	(0.479)	(4.615)	(3.284)
2020year	0.115	−0.391	−0.421	−0.257	−0.251	5.174	0.076
	(0.487)	(0.635)	(0.501)	(0.529)	(0.482)	(4.645)	(3.285)
Download speed (in logs)	0.059	0.049	0.019	0.036	0.015	0.477	−0.660
	(0.066)	(0.085)	(0.067)	(0.071)	(0.065)	(0.625)	(0.443)
2019year*Download speed	0.002	−0.031	0.006	−0.010	0.010	−0.038	−0.043
	(0.092)	(0.119)	(0.094)	(0.099)	(0.091)	(0.886)	(0.626)
2020year*Download speed	−0.017	0.022	0.036	0.018	0.019	0.238	0.710
	(0.093)	(0.120)	(0.095)	(0.100)	(0.091)	(0.894)	(0.626)
Constant	19.717***	19.136***	19.865***	20.417***	20.356***	42.752***	85.273***
	(0.345)	(0.450)	(0.355)	(0.375)	(0.342)	(3.254)	(2.325)
Observations	1,478	1,286	1,286	1,286	1,286	1,208	3,571
R^2	0.003	0.003	0.003	0.002	0.002	0.032	0.008
Adjusted R^2	−0.0003	−0.001	−0.001	−0.002	−0.002	0.028	0.006

Note: OLS robust standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01

Table 5: Granger-type causality tests for specification 2 (dummy)

	ACT scores						
	Composite	English	Math	Reading	Science	AP scores	Graduation rate
2019year	−0.186	−0.130	−0.139	−0.140	−0.151	−0.533	0.361
	(0.152)	(0.197)	(0.155)	(0.164)	(0.150)	(1.430)	(1.426)
2020year	−0.031	−0.311	−0.268*	−0.198	−0.175	6.745***	4.918***
	(0.152)	(0.197)	(0.156)	(0.165)	(0.150)	(1.441)	(1.424)
Low quality	−0.292	−0.351	−0.171	−0.253	−0.181	−0.656	2.409
	(0.202)	(0.265)	(0.209)	(0.221)	(0.202)	(1.909)	(1.508)
2019year*Low quality	0.118	0.166	0.107	0.196	0.116	0.650	0.097
	(0.285)	(0.374)	(0.295)	(0.312)	(0.284)	(2.708)	(2.132)
2020year*Low quality	0.218	0.123	0.113	0.112	0.074	−1.366	−2.018
	(0.285)	(0.374)	(0.295)	(0.312)	(0.284)	(2.718)	(2.132)
Constant	20.099***	19.483***	20.009***	20.672***	20.483***	45.332***	80.232***
	(0.107)	(0.161)	(0.110)	(0.116)	(0.106)	(1.007)	(1.008)
Observations	1,478	1,286	1,286	1,286	1,286	1,208	2,543
R^2	0.004	0.005	0.003	0.003	0.002	0.031	0.009
Adjusted R^2	0.0002	0.001	−0.001	−0.001	−0.002	0.027	0.007

Note: OLS robust standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01

treatment effects β_1 should equal 0. Therefore, when running the Granger-type causality tests, the hypothesis $H_1 : \beta_1 \neq 0$ is tested against the null hypothesis $H_0 : \beta_1 = 0$. Therefore, a lack of significance for β_1 in the Granger-type causality test would fail to reject the null hypothesis ($\beta_1 = 0$), which then provides evidence that there are no pre-treatment effects and further support for the viability of the parallel trends assumption.

The results of the Granger-type causality tests can be found in Table 4 for the first specification (Equation 3) and in Table 5 for the second specification (Equation 4). Looking at the coefficients on the interaction term between 2019 and the internet variable, there appears to be no pre-treatment effect for any of the dependent variables for either specification. Across all five ACT score variables, AP scores, and graduation rates, β_1 is not statistically different from 0. None of these coefficients are statistically significant at any standard significance level, indicating a failure to reject the null hypothesis that $\beta_1 \neq 0$. This would suggest that there is no pre-treatment effect and provides additional evidence that the parallel trends assumption is viable.

6 Results

In this section, I discuss the results of the two difference-in-differences specification for each of my education outcome variables. For the ACT data, I present regression results on the scores for each section of the exam as well as the composite score (both with and without Louisiana data) in one table for each specification. The results for Specification 1 can be found in Table 6 and the results for Specification 2 can be found in Table 7. The discussion of ACT results can be found in Section 6.1. Results for the AP scores and graduation rates are presented individually, with separate tables for each specification. The results for AP scores can be found in Table 8 (Specification 1) and Table 9 (Specification 2) with a discussion about them in Section 6.2. The results for graduation rates can be found in Table 10 (Specification 1) and 11 (Specification 2) with the discussion in Section 6.3.

All the results are summarized in regression tables with estimated coefficients presented and p-values obtained from wild bootstrapping in parentheses. I provide p-values instead of standard errors as the wild bootstrap method I use only outputs p-values. This is because the wild bootstrap method does not assume the underlying distribution follows a standard distribution (such as a normal distribution) making the calculation of standard errors impractical, especially in instances where the confidence intervals are asymmetric.

The coefficients of interest for these regressions are the DID estimators, the coefficients on the interaction terms. The interaction terms are labeled as *2020*Download speed* for when the internet variable is continuous (Specification 1) and as *2020*Low quality* when the internet variable is binned (Specification 2).

6.1 ACT Scores

Tables 6 and 7 present the DID results for the ACT scores. The regression for composite scores can be found in column 1. Column 2 also presents results for ACT composite scores, but this regression does not include data from Louisiana. I run an alternative regression without Louisiana data because Louisiana seems to have unusually positive results for composite scores between 2019 to 2020 compared to the other states and because I do not have ACT individual section results for Louisiana, thus removing this data makes column 2 more comparable to columns 3 to 6. The DID regression results for the English, Mathematics, Critical Reading, and Science sections (the individual sections of the ACT) can be found in columns 3, 4, 5, and 6, respectively.

Looking at the results of Specification 1 for the ACT scores, the DID estimator for the composite score without the Louisiana data (column 2) and the estimators for three of the four individual sections (columns 3-5; all except for Science) are positive. The composite scores without Louisiana, the English scores, and the reading scores all have the same coefficient (0.004) and the math scores have a larger coefficient (0.014). Because these results are positive, this would indicate that a one unit log increase in median internet download speed for a county would result in higher composite scores (no Louisiana), English scores, and reading scores by 0.004 and math scores by 0.014 in 2020 compared to previous years, holding all else equal. This would suggest that students in counties that have higher quality internet benefit more from the better internet in 2020. The signs of these coefficients are expected based on my hypotheses as they suggest that those with lower quality internet would have lower scores in 2020 than those with higher internet compared to previous years. However, while the sign of these coefficients matches my hypothesis, none of these results are statistically significant. These coefficients have p-values ranging from 0.69 (math) to 0.84 (science) which are very large and far from the traditionally acceptable significance thresholds of < 0.05 or < 0.01. This suggests that I cannot reject the null hypothesis that the coefficients are statistically different from 0.

For Specification 1 using ACT data, The only regressors with a negative DID estimator is composite scores with all the data (column 1), which has a coefficient of -0.037, and science scores (column 6), which has a coefficient of -0.002. These negative coefficients suggest that, in these regressions, the scores of students in a county that has lower quality internet, is higher, relative to previous years (lower for students from counties with higher quality internet). Specifically, the coefficients suggest that for each one log unit higher median download speed a county has, the composite test scores for that county's students was 0.016 points lower and the science test scores for that county's students was 0.002 points lower in 2020 compared to previous years, holding all else equal. These coefficients have the opposite signs of the expectations from my hypothesis as I anticipated that students from counties with higher quality internet would have achieved higher comparative test scores in 2020. However, like the DID estimators for the other ACT score regressions, neither the coefficient for composite scores with all data nor the coefficient for science scores is significant. The p-value for the composite scores is 0.22 and the p-value for science scores is 0.95 which are both larger than the critical values of 0.05 and 0.01 indicating a failure to reject the null hypothesis that these DID estimators are statistically different from 0.

Despite the lack of significance on the DID estimators, the coefficient on 2020 is significant at the 5% level for math scores (-0.25, $p = 0.04$). This coefficient represents the average change in ACT math scores in 2020 compared to previous years, holding all else constant. This coefficient suggests that, on average, students in my sample received 0.25 fewer points on the math section of the ACT in 2020 relative to previous years, holding all else equal. While this coefficient provides an interesting exploratory analysis because of its significance, no causal inference should be taken from it with regards to the onset of the pandemic or online learning. I did not formulate any hypotheses about this nor are my specifications designed to correctly detect this effect. There are likely omitted variables influencing this effect such as the difficulty of the exam.

Table 6: ACT scores results (continuous)

	Dependent variables:					
	Composite	Composite (no LA)	English	Math	Reading	Science
2020	0.296*	−0.180	−0.258*	−0.252**	−0.152	−0.088
	(0.082)	(0.134)	(0.077)	(0.039)	(0.268)	(0.464)
2020*Upload speed	−0.037	0.004	0.004	0.014	0.004	−0.002
	(0.222)	(0.871)	(0.881)	(0.525)	(0.877)	(0.951)
County-level Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,478	1,286	1,286	1,286	1,286	1,286
R^2	0.000	0.002	0.002	0.002	0.001	0.001

Note: p-values from wild bootstrapping in parentheses. *p<0.1; **p<0.05; ***p<0.01

When looking at the results for Specification 2 (binned internet variable) using ACT data, there are similar results. In this specification, all of the DID estimators are positive with the four subject sections and the composite scores without Louisiana data having quite small coefficients (< 0.1) and the coefficient on composite score with all data being 0.234. These results, however, are the opposite of what was expected from my hypothesis. Positive results on the DID estimator in this specification means that in the school year where schools shifted to online learning, the low-quality internet group would have higher ACT scores in all subjects and the composite score compared to their previous than the high-quality internet group would have, holding all else equal. These results would suggest that the shift to online learning would have less negatively affected the low-quality internet than the high-quality internet group. However, as with the previous specification, none of these results are statistically significant. The DID estimator for the composite score regression with all the data has the smallest p-value at 0.34, which is still far from acceptable thresholds

Table 7: ACT scores results (dummy)

	Dependent variables:					
	Composite	Composite (no LA)	English	Math	Reading	Science
2020	0.028	−0.171	−0.247	−0.201	−0.127	−0.101
	(0.865)	(0.215)	(0.162)	(0.143)	(0.381)	(0.470)
Low quality	−0.280**	−0.216	−0.283	−0.211	−0.171	−0.192
	(0.030)	(0.140)	(0.107)	(0.126)	(0.252)	(0.150)
2020*Low quality	0.234	0.038	0.041	0.061	0.011	0.021
	(0.335)	(0.869)	(0.900)	(0.796)	(0.956)	(0.928)
COVID-19 rate	−0.027***	−0.018***	−0.021***	−0.018***	−0.017***	−0.016***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Constant	20.217***	20.335***	19.557***	20.079***	20.713***	20.527***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	1,475	1,286	1,286	1,286	1,286	1,286
R^2	0.028	0.015	0.015	0.016	0.012	0.013
Adjusted R^2	0.025	0.012	0.012	0.013	0.009	0.010

Note: p-values from wild bootstrapping in parentheses. *p<0.1; **p<0.05; ***p<0.01

for significance. This means the null hypotheses cannot be rejected as the coefficients are not significantly different from 0.

However, one main coefficient in my analysis for ACT data using Specification 2 has significance. The coefficient on *Low quality* has a p-value of 0.03, so it is significant at the 5% level. The value of this coefficient, -0.28, indicates a negative relationship between being in the low-quality internet group and the ACT composite scores. Specifically, it indicates that, on average across the years, students from counties in the low-quality internet group score 0.28 points lower on the composite score, holding all else equal. Despite this, no causal inference should be taken from this result. My specification is not specifically designed to test this relationship so there could by omitted variable bias. Additionally, it is only significant when the Louisiana data is included and is not significant ($p = 0.14$) when the Louisiana data is removed, so this result should not be considered anything more than an interesting exploratory result.

Because of the lack of significance for the DID estimator in all six of these regressions for both specifications, I do not have any evidence to support my hypothesis that the shift to online learning impacted ACT scores for student from counties with lower quality internet more negatively than students from counties with higher quality internet.

6.2 AP Scores

The results for estimations using percentage of qualifying AP scores follows a similar pattern to the results for ACT scores. These results can be found in Table 8 for the results with the continuous internet variable (Specification 1) and in Table 9 for the results with the binned internet variable (Specification 2).

Looking at the results for Specification 1, the DID estimator is -0.102, which is negative, the opposite sign from what I expected based on my hypothesis. This coefficient indicates that students in a county with one log unit higher median speed internet are, on average, expected to receive a qualifying score on AP exams 0.10% less often in 2020 than previous years, holding all else equal. Thus, conversely, students from counties with worse internet quality are more likely to score more positively in 2020 compared to previous years, which is the opposite of my proposed hypothesis. Unfortunately, there is a lack of significance ($p = 0.689$) so the null hypothesis that the coefficient is significantly different from 0 cannot be rejected. Therefore, there is no evidence supporting my hypothesis.

Even with the lack of significance on the DID estimator, there is significance on the 2020 coefficient. This coefficient has a value of 6.20, indicating that, in general, AP scores were higher in 2020 compared to previous years. Specifically, on average, 6.2% more students received qualifying AP scores in 2020 than previous years, holding all else equal. This result is significant at the 1% level with a p-value of 0.00. While this result is interesting, no causal relationship should be inferred from it as my specifications were not designed to detect this effect. Instead, this offers some exploratory analysis that may be worth exploring in future research.

The results of Specification 2 for the percentage of qualifying AP scores are found in Table 9. Looking

Table 8: Percentage of qualifying AP scores results (continuous)

	Dependent variable:
	AP scores
2020	6.198***
	(0.000)
2020*Download speed	−0.102
	(0.689)
County-level Fixed Effects	Yes
Observations	1,206
R^2	0.030

Note: p-values from wild bootstrapping in parentheses. *p<0.1; **p<0.05; ***p<0.01

Table 9: Percentage of qualifying AP scores results (dummy)

	Dependent variable:
	AP scores
2020	7.220***
	(0.000)
Low quality	−0.025
	(0.994)
2020*Low quality	−2.080
	(0.330)
COVID-19 rate	−0.303***
	(0.001)
Constant	46.459***
	(0.000)
Observations	1,206
R^2	0.039
Adjusted R^2	0.036

Note: p-values from wild bootstrapping in parentheses. *p<0.1; **p<0.05; ***p<0.01

at the DID estimator in this regression, the sign again matches my expectations. Here, the β_2 coefficient is -2.08, suggesting that there is a negative effect for the low-quality internet group compared to the high-quality internet group. This suggests that the low-quality internet group performs less positively or more negatively on the 2020 AP exams compared to the exams from previous years when compared to the results of the high-quality internet group. Again, however, the results suffer from a lack of significance as the p-value on the DID estimator is 0.330, above the 0.05 threshold. Therefore, the null hypothesis ($\beta_2 = 0$) cannot be rejected and there is no evidence supporting my hypothesis.

Despite the lack of significance on the DID estimator, the coefficient on 2020 is significant (7.22, $p = 0.00$). This coefficient represents the mean change in the percentage of qualifying AP scores for the high-quality internet group in 2020. This coefficient suggests that counties with high-quality internet saw 7.22% more students with qualifying AP exam scores in 2020 than in previous years, holding all else equal. While this coefficient provides an interesting exploratory analysis because of its significance, no causal inference should be taken from it with regards to the onset of the pandemic or online learning. I did not formulate any hypotheses about this nor are my specifications designed to correctly detect this effect. There are likely omitted variables influencing this effect such as the difficulty of the exam. The AP exams in 2020 may have just been easier than previous years leading to a positive coefficient that is unrelated to COVID-19.

6.3 Graduation Rates

The results for the graduation rates regressions can be found in Table 10 with the continuous internet variable (Specification 1) and in Table 11 with the binned internet variable (Specification 2).

The results for graduation rates using Specification 1 provides the only significant result on my DID estimators. The DID estimator for this regression has a coefficient of 0.741, indicating a positive relationship, which matches with my hypothesis. This coefficient indicates that schools in counties with one log unit higher median download speed have, on average, a 0.74 higher graduation rate (out of 100.00) in 2020 compared to previous years, holding all else equal. This coefficient suggests that counties with higher-quality internet are comparatively better off in 2020 than counties with lower-quality internet. This indicates that the converse is true as well, like I propose in my hypothesis. This would suggest that for each one log unit lower median download speed a county has, the schools in that county are expected to have 0.74 lower graduation rate, holding all else equal. This result is significant at the 1% level with a p-value of 0.003. This provides evidence in support of my hypothesis (through sub-hypothesis H1c) that the shift to online learning had a more negative effect on education outcomes (specifically graduation rates) for students with lower quality internet compared to students with higher quality internet in 2020 relative to previous years.

Unlike the continuous specification, however, the regression for graduation rates using Specification 2 is insignificant. The sign on DID estimator for this regression (-2.05) matches my hypothesis. This coefficient suggests that, on average, schools in the low-quality internet group have a lower graduation rate by 2.05 percentage points in 2020 compared to the counterfactual represented by the high-quality internet group. However, this DID estimator lacks statistical significance ($p = 0.23$). Therefore, the null hypothesis that it is significantly different from 0 cannot be reject and it does not provide evidence to support my hypothesis.

While the DID estimator lacks significance, both variables that make up the interaction term have varying levels of significance with the coefficient on 2020 (4.74) significant at the 1% level and the coefficient on *Low quality* (2.39) significant at the 5% level. The coefficient on 2020 suggests that counties with high-quality

Table 10: Graduation rate results (continuous)

	Dependent variable:
	Graduation rate
2020	−0.032
	(0.972)
2020*Download speed	0.741***
	(0.003)
County-level Fixed Effects	Yes
Observations	2,543
R^2	0.006

Note: p-values from wild bootstrapping in parentheses. *p<0.1; **p<0.05; ***p<0.01

Table 11: Graduation rate results (dummy)

	Dependent variable:
	Graduation rate
2020	4.737***
	(0.000)
Low quality	2.389**
	(0.032)
2020*Low quality	−2.049
	(0.227)
COVID-19 rate	0.113**
	(0.021)
Constant	79.932***
	(0.000)
Observations	2,540
R^2	0.010
Adjusted R^2	0.009

Note: p-values from wild bootstrapping in parentheses. *p<0.1; **p<0.05; ***p<0.01

internet saw 4.74% higher graduation rates in 2020 than in previous years, holding all else equal. The coefficient on *Low quality*, which represents the baseline difference between the low-quality internet and high-quality internet groups, suggests that the low-quality internet group had a 2.39% higher graduation rate prior to 2020 than the high-quality internet group, holding all else equal. However, no causal inference in regards to the shift to online learning or the the onset of the COVID-19 pandemic should be taken from these coefficients as the specifications were not designed to unbiasedly measure these effects and there are likely omitted variables that are influencing these results. The significance of these coefficients just offers some interesting exploratory analysis.

7 Robustness Checks

To test the soundness of my main results, I run a series of two robustness checks using alternative specifications. Both of my robustness specifications use modified versions of Specification 1 (Equation 1) and Specification 2 (Equation 2) from my main analysis but with the alternative internet variables of average download speed and median upload speed. Like most of my primary results, most of my robustness specifications find similarly non-significant results on the DID estimator. For the continuous internet variable specification with graduation rates, my robustness checks find statistically significant results on the DID estimator, offering additional support in favor of my hypothesis. While I briefly discuss the specifications and results of the robustness checks in this section, the tables of the full outputs can be found in Appendix B.

Average Download Speeds

The first set of robustness checks use average download speed instead of median download speed. Using similar specifications to the main analysis, these regressions have one specification using the logged version of the continuous internet variable and one specification using the binned internet variable. Average download speed is binned using the same conditions that median download speed is binned, specifically that the low-quality bin consists of counties with average download speed $\leq$ 50 Mbps and the high-quality bin consists of counties with average download speed > 50 Mbps. The continuous specification, using county-level fixed effects, can be written as:

$$edu_{yc} = \beta_1 \cdot 2020year_{yc} + \beta_2 \cdot (Downloadspeed_c \times 2020year_{yc}) + a_c + \epsilon_{yc} \tag{5}$$

where $Downloadspeed_c \times 2020year_{yc}$ is the DID estimator that is the interaction between 2020 and average download speed and the other variables are the same as the primary regression. The binned specification can be written as:

$$edu_{yc} = \alpha_1 + \alpha_2 \cdot 2020year_{yc} + \beta_1 \cdot Lowquality_c + \beta_2 \cdot (Lowquality_c \times 2020year_{yc}) + \gamma_1 \cdot COVIDrate_{yc} + \epsilon_{yc} \tag{6}$$

where $Lowquality$ equals 1 if average max advertised download speed is $\leq$ 50 and 0 otherwise and the other variables are the same as the primary regression. Both of these regressions use unclustered wild bootstrapping for the standard errors and the continuous specification uses county-level fixed effects.

Looking at the results for the robustness test using average download speed, I find they match my primary results as none of the DID estimators are statistically significant except for the interaction in the regression of graduation rate using the continuous internet variable. For ACT scores, p-values range from 0.68 to 0.97 in the continuous internet variable specification (Table B.1) and from 0.60 to 0.91 in the the binned internet variable specification (Table B.2). Similar results are found in the AP scores (Table B.3 and Table B.4) with p-values of 0.75 for the continuous internet variable specification and 0.70 for the binned internet variable specification. None of these results are below the critical value of $p < 0.05$ thus none of them are significant. This helps provide support for the insignificance of the main results for ACT scores and AP test scores.

For the continuous internet variable specification of graduation rates, this robustness check adds support for the significance of my main results as this regression also finds significance for the DID estimator (Table B.5). While not significant at the 1% level like the main results, the results in this robustness test find it significant at the 5% level ($p = 0.05$). Additionally, the coefficient of 0.86 in this robustness check has the same sign and similar magnitude to the main result of 0.74 (Table 10). However, like the main results, the binned internet variable specification of graduation results finds insignificant results (Table B.6). This specification finds a p-value of 0.69 on the DID estimator which is above the 5% significance level.

Median Upload Speeds

The second set of robustness checks use median upload speed instead of median download speed. Using similar specifications to the main analysis, these regressions have one specification using the logged version of the continuous internet variable and one specification using the binned internet variable. Median upload speed is binned in a method where the low-quality bin consists of counties with median upload speed ≤ 5 Mbps and the high-quality bin consists of counties with median upload speed > 5 Mbps. The continuous specification, using county-level fixed effects, can be written as:

$$edu_{yc} = \beta_1 \cdot 2020year_{yc} + \beta_2 \cdot (Uploadspeed_c \times 2020year_{yc}) + a_c + \epsilon_{yc} \tag{7}$$

where $Uploadspeed_c \times 2020year_{yc}$ is the DID estimator that is the interaction between 2020 and median upload speed and the other variables are the same as the primary regression. The binned specification can be written as:

$$edu_{yc} = \alpha_1 + \alpha_2 \cdot 2020year_{yc} + \beta_1 \cdot Lowquality_c + \beta_2 \cdot (Lowquality_c \times 2020year_{yc}) + \gamma_1 \cdot COVIDrate_{yc} + \epsilon_{yc} \tag{8}$$

where $Lowquality$ equals 1 if median max advertised upload speed is ≤ 5 and 0 otherwise and the other variables are the same as the primary regression. Both of these regressions use unclustered wild bootstrapping for the standard errors and the continuous specification uses county-level fixed effects.

With the robustness check using median upload speed, I again find similarly insignificant results for most of the education variables in both specifications. The results for ACT scores has no significant DID estimators in either the continuous internet variable specification (Table B.7) or the binned internet variable specification (Table B.8) with p-values ranging from 0.46 to 0.87 and 0.61 to 0.81, respectively. The results for AP scores (Table B.9 and Table B.10) have equally insignificant results with p-values of 0.16 and 0.96 on the DID estimators, respectively. Since none of these results are below the critical value of $p < 0.05$ they are not significant, which helps provide support for the insignificance of the main results for ACT scores and AP test scores.

The results for graduation rates in this second robustness check, however, again find significance for the DID estimator in the continuous internet variable specification. While upload speeds are a different measure than download speeds, thus the magnitudes cannot be compared, they are correlated due to the fact that counties with higher download speeds tend to have higher upload speeds ($\rho = 0.867, p = 0.000$).[5] Additionally, upload speed is another measure of internet quality, so better upload speeds would suggest better internet quality. Looking at the DID estimator in this robustness check, the coefficient, 0.57 (Table B.11), is positive meaning it has the same sign as the coefficient of the main result, 0.74 (Table 10). This result is also significant, although unlike the main result, it is only significant at the 5% level and not the 1% level. The graduation rate results for the DID estimator in the binned internet variable is not significant at the 5% level with a p-value of 0.08 (Table B.12).

Both of these robustness checks add support for the main results on the continuous graduation rates indicating that it is less likely to be a false positive. This also gives greater evidence for my hypothesis that better internet quality more negatively or less positively affects the education outcomes, specifically with regards to graduation rates, of students with lower-quality internet compared to students with high-quality internet in 2020 relatively to previous years.

8 Discussion

Within my results, I find some evidence in support of my hypothesis that shifting to online classes due to the COVID-19 pandemic affects students with low-quality internet disproportionately more negatively or less positively than students with high-quality internet. Looking at the graduation rates variable under Specification 1 with the continuous internet variable, I have a robust statistically significant result indicating that the better internet quality is for a county, the comparatively higher graduation rate schools in that county have in 2020 relative to previous years. This would suggest that, with the shift to online learning, that students with better internet quality were more likely to graduate, relative to previous years, than students with worse internet quality. This also suggests, in alignment with how my hypothesis is structured, that students from counties with low-quality internet would be comparatively less likely to graduate in 2020 relative to the graduation rates in previous years.

Not only does this support my hypothesis, but this makes empirical sense as well. Students with lower quality internet are more likely to have issues with attending online classes and completing internet-intensive tasks and projects. This could cause them to struggle in classes and receive lower grades than they otherwise might have if the classes were in-person. If this is the case, students that were borderline passing may have been unable to overcome the internet quality deficiency enough to pass critical classes. Another possibility

[5] ρ is the Spearman rank-order correlation coefficient.

is that students with low-quality internet get frustrated with the difficulties caused by their poorer quality internet and the difficulty in attending and managing courses. This frustration could translate into even poorer grades. These struggles with the internet could also translate into students being unable to properly attend class so they decide to defer graduation until classes resume in-person learning. All of these are reasonable possibilities for why students with lower quality internet may not graduate, thus leading to comparatively lower graduation rates for these counties in 2020.

Throughout the rest of my results, there is a distinct lack of statistical significance on the coefficients of interest both within my primary results and my robustness checks. For all the ACT scores data and the AP test scores data across both of the specifications, none of the DID estimators are statistically different from 0. There could be a few reasons for why this has happened.

First, the true effects for these variables could actually be 0, indicating that having different internet quality does not effect the education outcomes of students measured here (ACT scores and AP test scores) after a shift to online learning. This would suggest that students from counties with low-quality internet are not worse off than students from counties with high-quality internet. This is certainly a possible explanation. The majority of the students in counties from low-quality internet may have had enough internet to succeed at distance learning. Despite the 50 Mbps for online learning for students I proposed earlier (Section 4.1), there is the possibility that lower internet speeds are suitable enough for distance learning success, at least in its relation to taking standardized tests. Another possibility is that these families adapted quickly and successfully to sharing their lower speed internet with each other ensuring that the family members could succeed at their school and work despite the lower-quality internet. Another potential reason is that these students found alternative solutions for their internet problems. One possibility is that school or community programs could have been developed to help students with lower internet access. For example, communities may have provided students who have low-quality internet with temporary devices that had better internet access or schools may have developed lessons designed more for "at-home" study that relied less on the internet.

Additionally, a reason the actual effect may be zero is related specifically to my internet data. I use fixed broadband internet which is still the primary internet form used in households as it is the source of WiFi. However, with the improved sophistication of smartphones and tablets and dropping data prices, using mobile devices has increasingly become a viable alternative to using WiFi. This means that, while households have low-quality broadband internet, they may have fast enough mobile internet to compensate, allowing students to attend lectures and succeed in online classes without higher quality internet. If enough of the students from counties with low-quality fixed broadband internet can compensate for it using mobile internet, the effect difference between groups when shifting to online learning would be negligible. However, one issue with these possible reasons is the robust significance on graduation rates under the continuous internet variable specification. It seems unlikely that students from counties with low-quality internet would be successful in adjusting or compensating for their internet deficiency in regard to standardized testing, but not in regard to graduation. Instead, assuming the statistically significant result on graduation rates is not a false positive, an alternative reason for the insignificance with ACT and AP test scores seems more plausible.

One alternative reason for why my results would be statistically insignificant even if the shift to online learning during the COVID-19 pandemic would affect students in counties with low-quality internet disproportionately more is that the effect size is small and my study is not sufficiently powered to detect the small effects. The limited number of counties that data was available I reduces my sample size which reduces the power within this paper. Additionally, because of the nature of ACT scores and AP test scores, it is unlikely for counties to see large changes from year-to-year even in times with disruptive events like COVID-19. This suggests that any effects, while present, may have been too small for this paper to detect with any significance.

Another alternative is that, while the shift to distance learning does have a disproportionate affect on students from counties with low-quality internet, the effect would not be seen in the 2020 data. Because I look at standardized test scores in 2020, which occurred only a few months after the onset of the COVID-19 pandemic and the shift to online classes, there may have not been a long enough amount of time to see a noticeable effect. The majority of the school year would have taken place in person under normal conditions so any effect from online learning may have been minimal, despite happening at a pivotal time in the preparation process. For example, the standardized tests may have consisted of content based on material that was primarily learned prior to the shift which would reduce the effects from shifting to online learning. Because of these factors, looking at a period of time where the shift to online learning was longer, such as looking at 2021 data, could allow for detection of the actual effects. This could be an interesting avenue for future researchers to look into, however, other issues would need to be addressed while doing this analysis, such as compensating for changes to college applications structure where they no longer require college readiness tests and as well as adjusting for changes to the overall structure of the ACT and AP tests.

9 Limitations

While I tried to make the results of this paper as robust as possible, there are several limitations to the analysis. The first limitation with my paper is the data, particularly the lack of data available. I discuss the main issues with data collection in Section 4. Because there was no centrally available dataset that I had access to for any of the education outcomes variables, I had to collect the data from the individual state websites, which each offered different amounts of data and data in different formats. Many states did not even provide data for the most recent year (or in some cases several most recent years). The limited number of states that provide data (seven for the ACT and AP test scores and 16 for graduation rates) leads to two potential issues in my paper. One is a reduced sample size. With counties from all fifty states, there would be much greater power which would allow for better detection of small effect sizes. Additionally, because the restricted sample is only specific states, it does not provide as good of a cross-section of the United States' population which leads to a loss of generalizability. While my results can provide robust analysis for the states within my sample, it lacks in the external validity needed to make broader claims about the entirety of the United States.

Another potential limitation is omitted variable bias. Most analysis has the potential to suffer from omitted variable bias as it is difficult to determine and find data for all the variables that potentially affect an analysis. While the construction of the difference-in-differences analysis and the use of fixed effects in Specification 1 helped to reduce potential omitted variable bias, lack of data prevented me from including all the time-variant control variables I identified as potential omitted variables, specifically changes in county-level education spending. For the states included in this analysis, I was unable to find education expenditure data for all of them at a county level and, even for those I found data for, it was not structured in a comparable form. Due to this, I have not included this control in my analysis, and while I do not believe it would have a large impact on the results, there is still the potential that it is causing omitted variable bias and this should be considered when interpreting the results in this paper. Another potential omitted variable has to do with the ability of a household to cope with the shift to online. Students from households with better socioeconomic conditions may be better prepared to handle to classes online compared to students from lower socioeconomic households, independent of internet. For example, students from higher socioeconomic households may have parents who can easily work from home and help their children with schoolwork, while parents in lower socioeconomic households may have less time to help. Since socioeconomic conditions are correlated with internet (Gorski, 2005;Zhao et al., 2010), this could lead to potential omitted variable bias. However, I do not think this would be too problematic since at least some of this effect would be present prior to 2020, and any added effect during the pandemic is likely minimal. However, the strong correlation between socioeconomic conditions and internet quality could mean that part of the effect I find in this paper for graduation rates is actually measuring differences in socioeconomic and not just internet quality, which would alter the interpretation of the results. While I do not look specifically at separating the effect of socioeconomic conditions from internet inequality in this paper, this could be a good avenue for future research.

A final limitation is in the underlying assumptions of DID analysis, particularly the parallel trends assumptions. As I have done in this paper, the viability of the parallel trends assumption can be argued using various methods such as visual inspection of pre-trends and testing for pre-treatment effects. While the three-period data I use in this paper (two pre-treatment periods and one post-treatment period) is adequate to run such an analysis, it is generally preferable to have more time periods so long-term pre-trends can be analyzed and compared. However, limitations on the data available restricted the number of pre-treatment periods that could be used. Therefore, while I believe I have provided adequate argumentation for the viability of the parallel trends assumption, more pre-treatment periods would have allowed for stronger argumentation and more compelling evidence on the viability of the parallel trends assumption.

10 Conclusion

The "Digital Divide", the gap in access to and knowledge of technology between groups, arose as an issue with the rapid technological innovations of the late 20th Century, and it continues to persist as a problem into the modern day. While falling costs of technology and government backed initiatives are helping to reduce the gap in technology inequality, the divide is still far from being completely bridged. Researchers have become especially concerned with the digital divide due to the effect it has on commonly marginalized groups. Researchers have found that minorities, women, rural areas, and lower socioeconomic groups tend to be disproportionately affected by this technological divide (Gorski, 2005; Strover, 2001; Zhao et al., 2010). One area where this inequity in technology seems particularly pronounced is education. Many researchers have found that a lack of access to technology or a lack of knowledge in using that technology has led to less scholastic success (Attewell and Battle, 1999; Buzzetto-Hollywood et al., 2018; Hargittai, 2002). Lower test scores (Aydin, 2021; Huang and Russell, 2006), lower enrollment (Fairlie, 2005), and increased dropout rates (Mbunge et al., 2020) are all academic areas where researchers have found an effect from the digital divide.

Through this paper, I have attempted to expand this literature on the digital divide's effect on education by looking for additional evidence about the effect of internet inequality on education outcomes. Within this paper, I look beyond the often used dichotomy of "haves" and "have-nots" of internet inequity, and instead look at the quality of the internet. Specifically I focus on students within the United States and set out to answer the question: *How does the shift to online distance learning affect the education of American students with differing internet quality at home?* For my analysis, I exploit the natural phenomenon of the COVID-19 global pandemic, which forced most schools in the United States to shut down and to shift to online distance learning in late March of 2020. Using difference-in-differences specifications with both continuous and binned internet quality variables, I test my hypothesis that students with low-quality internet were disproportionately affected by the shift to online learning in 2020 relative to previous years and compared to students with high-quality internet. To do this, I look at three education outcome variables, ACT scores, AP test scores, and graduation rates, and I use county-level data for the selection of states that provide 2020 data as well as data for the 2018 and 2019 school years.

From my results, I find some evidence to support my hypothesis through the education outcome variable of graduation rates. Under the specification that utilizes the continuous internet variable, I get statistically significant results in the hypothesized direction in my main results and in both of my robustness checks. This suggests that students with low-quality internet are disproportionately affected by the shift to online learning in 2020 relative to the previous years, at least in regard to graduation rates. I expect this effect comes from the inability of students with low-quality internet to properly attend, manage, and perform in academic settings due to internet connectivity and speed issues, especially compared to the students with high-quality internet. This would then affect education outcomes like graduation rates as students may be receiving poorer grades, dropping out from school, or deferring graduation until after classes return to in-person. While the results for graduation rates have significance, the results of my other education outcome variables are not significant.

The significant effect found with graduation rates contributes to the literature by adding additional evidence of the effect of the digital divide on education outcomes as well as providing evidence that the quality of access to the internet has an effect. Additionally, it provides some preliminary evidence into the effect of shifting to online classes during the COVID-19 pandemic in 2020, an area lacking in research due to its recency. However, there are many avenues that future researchers could look into. One limitations with my research was the limited data I had access to in regards to my education outcome variables. It could be interesting to look at the full sample of states once the data becomes available and see if the results in this paper match the results from the the full sample of states. Another potential limitation I have is the limited time between the onset of the pandemic and the results of the education variables. Looking at a sample where there was a longer period of online learning before these standardized test scores and graduation could show more pronounced effects and would be interesting to look into. A final avenue of future research is to use internet speeds or other measures of internet quality to look at other effects of inequity in internet access. As far as I can tell, this is the first paper to utilize internet speed as a measure of internet quality to look at the effects of the digital divide on education. It could be interesting to use these metrics to study other effects of this divide as well.

References

Abadie, A., Athey, S., Imbens, G. W., & Wooldridge, J. (2017). *When should you adjust standard errors for clustering?* [Working Paper No. 24003]. National Bureau Of Economic Research. http://www.nber.org/papers/w24003

Acemoglu, D., Autor, D. H., & Lyle, D. (2004). Women, war, and wages: The effect of female labor supply on the wage structure at midcentury. *Journal of Political Economy*, *112*(3), 497–551. https://doi.org/10.1086/383100

ACT, Inc. (2018). *The ACT profile report - national: Graduating class of 2018*.

ACT, Inc. (2020). *The ACT profile report - national: Graduating class of 2020*.

Angrist, J., & Lavy, V. (2002). New evidence on classroom computers and pupil learning. *The Economic Journal*, *112*, 735–765. https://doi.org/10.1111/1468-0297.00068

Angrist, J., & Pischke, J. (2010). The credibility revolution in empirical economics: How better research design is taking the con out of econometrics. *Journal of Economic Perspectives*, *24*(2), 3–30. https://doi.org/10.1257/jep.24.2.3

Attewell, P. (2001). Comment: The first and second digital divides. *Sociology of Education*, *74*(3), 252–259. https://doi.org/10.2307/2673277

Attewell, P., & Battle, J. (1999). Home computers and school performance. *The Information Society*, *15*(1), 1–10. https://doi.org/10.1080/019722499128628

Aydin, M. (2021). Does the digital divide matter? Factors and conditions that promote ICT literacy. *Telematics and Informatics*, *58*, 101536. https://doi.org/10.1016/j.tele.2020.101536

Bell, K. (2021). Senators ask the fcc to change the definition of high-speed broadband. *Yahoo Finance*. https://finance.yahoo.com/news/senators-fcc-change-definition-high-speed-broadband-222150947.html

Bertrand, M., Duflo, E., & Mullainathan, S. (2004). Bootstrap-based improvements for inference with clustered errors. *The Quarterly Journal of Economics*, *119*(1), 249–275. https://doi.org/10.1162/003355304772839588

Bettinger, E. P., Fox, L., Loeb, S., & Taylor, E. S. (2017). Virtual classrooms: How online college courses affect student success. *American Economic Review*, *107*(9), 2855–2875. https://doi.org/10.1257/aer.20151193

Bolden, K. (2021). Internet speed: How much do you really need? *CNET*. https://www.cnet.com/home/internet/how-much-internet-speed-do-you-really-need/

Bucy, E. P. (2000). Social access to the internet. *Harvard International Journal of Press/Politics*, *5*(1), 50–61. https://doi.org/10.1177/1081180X00005001005

Buzzetto-Hollywood, N., Wang, H., Elobeid, M., & Elobaid, M. (2018). Addressing information literacy and the digital divide in higher education. *Interdisciplinary Journal of e-Skills and Lifelong Learning*, *14*, 77–93. https://doi.org/10.28945/4029

Cacault, M. P., Hildebrand, C., Laurent-Lucchetti, J., & Pellizzari, M. (2021). Distance learning in higher education: Evidence from a randomized experiment. *Journal of the European Economic Association*. https://doi.org/10.1093/jeea/jvaa060

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics*, *90*(3), 414–427. https://doi.org/10.1162/rest.90.3.414

Census Bureau. (2019). *Glossary*. https://www.census.gov/programs-surveys/geography/about/glossary.html#par_textimage_5

Chin, M. (2020). Students are failing ap tests because the college board can't handle iphone photos. *The Verge*. https://www.theverge.com/2020/5/20/21262302/ap-test-fail-iphone-photos-glitch-email-college-board-jpeg-heic

College Board. (2018). *AP cohort data report: Graduating class of 2018*.

College Board. (2019). *AP cohort data report: Graduating class of 2019*.

College Board. (2020a). *AP cohort data report: Graduating class of 2020*.

College Board. (2020b). Past AP exam dates. https://apstudents.collegeboard.org/about-ap-exams/past-exam-dates

Davidson, R., & Flachair, E. (2008). The wild bootstrap, tamed at last. *Journal of Econometrics*, *146*, 162–169. https://doi.org/10.1016/j.jeconom.2008.08.003

DiMaggio, P., & Hargittai, E. (2001). *From the 'digital divide' to 'digital inequality': Studying internet use as penetration increases* [Working Paper No. 15]. Center for Arts; Cultural Policy Studies, Princeton University.

Donald, S. G., & Lang, K. (2007). Inference with difference-in-differences and other panel data. *The Review of Economics and Statistics*, *89*(2), 221–233. https://doi.org/10.1162/rest.89.2.221

EducationWeek. (2020). Map: Coronavirus and school closures in 2019-2020. https://www.edweek.org/leadership/map-coronavirus-and-school-closures-in-2019-2020/2020/03

Fairlie, R. W. (2005). The effects of home computers on school enrollment. *Economics of Education Review*, *24*, 533–547. https://doi.org/10.1016/j.econedurev.2004.08.008

FCC. (2019). *Glossary of terms used in FCC Form 477 instructions*. https://us-fcc.app.box.com/v/Form477Glossary

FCC. (2020). *Household broadband guide*. fcc.gov/consumers/guides/household-broadband-guide

Figlio, D., Rush, M., & Yin, L. (2013). Is it live or is it internet? Experimental estimates of the effects of online instruction on student learning. *Journal of Labor Economics*, *31*(4), 763–784. https://doi.org/10.1086/669930

Gorski, P. (2005). Education equity and the digital divide. *Association for the Advancement of Computing In Education Review (formerly AACE Journal)*, *13*(1), 3–45.

Hargittai, E. (2002). Second-level digital divide: Differences in people's online skills. *First Monday*, *7*(4), 3–9.

Hargittai, E. (2005). Survey measures of web-oriented digital literacy. *Social Science Computer Review*, *23*(3), 371–379. https://doi.org/10.1177/0894439305275911

Haughton, J., & Kelly, A. (2015). Student performance in an introductory business statistics course: Does delivery mode matter? *Journal of Education for Business*, *90*(1), 31–43. https://doi.org/10.1080/08832323.2014.968518

Heerwegh, D., De Wit, K., & Verhoeven, J. C. (2016). Exploring the self-reported ICT skill levels of undergraduate science students. *Journal of Information Technology Education: Research*, *14*, 19–47.

Hess, A. J. (2020). Online ap exams start on monday—here's what students and parents need to know. *CNBC News*. https://www.cnbc.com/2020/05/07/online-ap-exams-start-on-mondayheres-what-students-need-to-know.html

Horaczek, S. (2020). Here's how much internet bandwidth you actually need to work from home. *Popular Science*. https://www.popsci.com/story/technology/work-from-home-broadband-connection-internet-fcc/

Huang, J., & Russell, S. (2006). The digital divide and academic achievement. *The Electronic Library*, *24*(2), 160–173. https://doi.org/10.1108/02640470610660350

Joyce, T. J., Crockett, S., Jaeger, D. A., Altindag, O., & O'Connell, S. D. (2015). Does classroom time matter? *Economics of Education Review*, *46*, 64–77. https://doi.org/10.1016/j.econedurev.2015.02.007

Kahn-Lang, A., & Lang, K. (2020). The promise and pitfalls of differences-in- differences: Reflections on 16 and Pregnant and other applications. *Journal of Business & Economic Statistics*, *38*(3), 613–620. https://doi.org/10.1080/07350015.2018.1546591

Katz, J. E., Rice, R. E., & Aspden, P. (2001). The internet, 1995-2000: Access, civic involvement, and social interaction. *American behavioral scientist*, *45*(3), 405–419. https://doi.org/10.1177/0002764201045003004

Kircher, M. M. (2020). Students think the college board is running a reddit sting to catch ap test cheaters. *Vulture*. https://www.vulture.com/2020/05/college-board-fake-reddit-account-ap-test-cheaters.html

Lewin, T. (2014). A new sat aims to realign with schoolwork. *The New York Times*. https://www.nytimes.com/2014/03/06/education/major-changes-in-sat-announced-by-college-board.html

Looker, E. D., & Thiessen, V. (2003). *The digital divide in Canadian schools: Factors affecting student access to and use of information technology*. Statistics, Canada, Ottawa.

MacKinnon, J. G. (2006). Bootstrap methods in econometrics. *The Economic Record*, *82*, S2–S18. https://doi.org/10.1111/j.1475-4932.2006.00328.x

MacKinnon, J. G. (2009). Bootstrap hypothesis testing. In D. A. Belsley & E. J. Kontoghiorghes (Eds.), *Handbook of Computational Econometrics* (pp. 183–213). Wiley Online Library. https://doi.org/10.1002/9780470748916.ch6

Mbunge, E., Fashoto, S. G., Akinnuwesi, B. A., Gurajena, C., Metfula, A. S., & Mashwama, P. (2020). *Covid-19 pandemic in higher education: Critical role of emerging technologies in Zimbabwe* [SSRN 3743246].

Mbunge, E., Fashoto, S. G., & Olaomi, J. (2021). *Covid-19 and online learning: Factors influencing students' academic performance in first-year computer programming courses in higher education* [SSRN 3757988].

Merkus, E., & Schafmeister, F. (2021). The role of in-person tutorials in higher education. *Economics Letters*, *201*(109801). https://doi.org/10.1016/j.econlet.2021.109801

Mogey, N., & Fluck, A. (2015). Factors influencing student preference when comparing handwriting and typing for essay style examinations. *British Journal of Educational Technology*, *46*(4), 793–802. https://doi.org/10.1111/bjet.12171

National Center for Education Statistics. (2020). Children's internet access at home [NCES 2020-144]. *The Condition of Education 2020* (pp. 10–14). U.S. Department of Education.

Peña-Lopez, I. (2010). From laptops to competences: Bridging the digital divide in education. *International Journal of Educational Technology in Higher Education (formerly RUSC. Universities and Knowledge Society Journal)*, *7*(1), 21–32.

Pierce, J. (2019). Digital divide. *The International Encyclopedia of Media Literacy*, 1–8. https://doi.org/10.1002/9781118978238.ieml0052

Richards, E., West, S., & Altavena, L. (2020). Amid coronavirus, AP exams went online and had tech problems. College Board says it's investigating. *USA Today*. https://eu.usatoday.com/story/news/education/2020/05/15/coronavirus-ap-exam-2020-college-board-troubleshooting/5194639002/

Riddlesden, D., & Singleton, A. D. (2014). Broadband speed equity: A new digital divide? *Applied Geography*, *52*, 25–33. https://doi.org/10.1016/j.apgeog.2014.04.008

Roodman, D., MacKinnon, J. G., Nielsen, M. Ø., & Webb, M. D. (2019). The wild bootstrap, tamed at last. *The Stata Journal*, *19*(1), 4–60. https://doi.org/10.1177/1536867X19830877

Roth, J. (2018). *Should we adjust for the test for pre-trends in difference-in-difference designs?* [No. 1804.01208]. arXiv.org. https://ideas.repec.org/p/arx/papers/1804.01208.html

Snouwaert, J. (2020). Nearly 10,000 students ran into issues submitting their ap exams because of technical glitches. *Business Insider*. https://www.businessinsider.com/students-experience-issues-submitting-online-ap-exams-2020-5?r=US&IR=T

Speedtest. (2021). Speedtest global index: Global speeds march 2021. https://www.speedtest.net/global-index

Strover, S. (2001). Rural internet connectivity. *Telecommunications Policy*, *25*, 331–347. https://doi.org/10.1016/S0308-5961(01)00008-8

Subramony, D. P. (2014). Revisiting the digital divide in the context of a 'flattening' world. *Educational Technology*, *54*(2), 3–9.

The Princeton Review. (2020). ACT test dates. https://www.princetonreview.com/college/act-test-dates

Warschauer, M., Knobel, M., & Stone, L. (2004). Technology and equity in schooling: Deconstructing the digital divide. *Educational Policy*, *18*(4), 562–588. https://doi.org/10.1177/0895904804266469

Wei, L., & Hindman, D. B. (2011). Does the digital divide matter more? comparing the effects of new media and old media use on the education-based knowledge gap. *Mass Communication and Society*, *14*(2), 216–235. https://doi.org/10.1080/15205431003642707

West, S., & Johnson, A. (2020). 'overwhelmed with the entire process': 2020 AP exams are online, thanks to coronavirus. *USA Today*. https://eu.usatoday.com/story/news/nation/2020/05/11/2020-ap-test-changes-coronavirus-online-essay-no-multiple-choice-rules-how-long-exam-study/3107814001/

Wing, C., Simon, K., & Bello-Gomez, R. A. (2018). Designing difference in difference studies: Best practices for public health policy research. *Annual Review of Public Health*, *39*, 453–469. https://doi.org/10.1146/annurev-publhealth-040617-013507

Wooldridge, J. M. (2012). Advanced panel data methods. *Introductory Econometrics: A Modern Approach 5th Edition* (pp. 484–511). South-Western Cengage Learning.

Wu, C. F. J. (1986). Jackknife, bootstrap and other resampling methods in regression analysis. *The Annals of Statistics*, *14*(4), 1261–1295. https://doi.org/10.1214/aos/1176350142

Zhang, M., Trussel, R. P., Tillman, D. A., & An, S. A. (2015). Tracking the rise of web information needs for mobile education and an emerging trend of digital divide. *Computers in the Schools*, *32*, 83–104. https://doi.org/10.1080/07380569.2015.1030531

Zhao, L., Lu, Y., Huang, W., & Wang, Q. (2010). Internet inequality: The relationship between high school students' internet use in different locations and their internet self-efficacy. *Computers & Education*, *55*, 1405–1423. https://doi.org/10.1016/j.compedu.2010.05.010

Zoom Help Center. (2021). System requirements for Windows, macOS, and Linux. https://support.zoom.us/hc/en-us/articles/201362023-System-requirements-for-Windows-macOS-and-Linux

Appendix A Summary Statistics

Table A.1: Internet download speed summary statistics by state

State	Minimum	Median	Mean	Max	State	Minimum	Median	Mean	Max
Alabama	2	100	398.2	1,000	Montana	25	100	352.3	1,000
Alaska	35	100	435.5	1,000	Nebraska	35	100	296.3	1,000
Arizona	35	100	433.5	1,000	Nevada	35	100	446.5	1,000
Arkansas	35	100	370.7	1,000	New Hampshire	35	100	282	1,000
California	2	100	358.3	1,000	New Jersey	35	250	388.3	1,000
Colorado	2	115	439.9	1,000	New Mexico	35	100	467	1,000
Connecticut	50	618.5	570.1	1,000	New York	35	100	369	1,000
Delaware	35	940	658.3	1,000	North Carolina	35	100	378.6	1,000
Florida	35	100	411	1,000	North Dakota	35	100	386.2	1,000
Georgia	35	100	368.7	1,000	Ohio	35	100	373.6	1,000
Hawaii	35	940	590	1,000	Oklahoma	2	100	474.1	1,000
Idaho	35	100	398.5	1,000	Oregon	35	325	512.4	1,000
Illinois	35	100	302.7	1,000	Pennsylvania	25	100	307	1,000
Indiana	35	225	503.7	1,000	Rhode Island	35	100	242	940
Iowa	2	100	416.1	1,000	South Carolina	35	100	376.7	1,000
Kansas	4	100	390.2	1,000	South Dakota	35	100	357.8	1,000
Kentucky	35	100	321.9	1,000	Tennessee	35	100	415.3	1,000
Louisiana	30	100	288.1	1,000	Texas	2	100	439.2	1,000
Maine	35	100	317.3	1,000	Utah	35	100	400	1,000
Maryland	35	175	400.7	1,000	Vermont	35	107.5	408	1,000
Massachusetts	35	100	397.6	1,000	Virginia	25	100	312.7	1,000
Michigan	35	100	405	1,000	Washington	2	100	372.8	1,000
Minnesota	35	100	325.8	1,000	West Virginia	35	100	346.6	1,000
Mississippi	35	100	362.5	1,000	Wisconsin	35	100	375.1	1,000
Missouri	2	100	348.6	1,000	Wyoming	35	100	382.2	1,000

Appendix B Robustness Checks Tables

Table B.1: ACT scores results with average internet download speed (continuous)

	Dependent variables:					
	Composite	Composite (no LA)	English	Math	Reading	Science
2020	0.181	−0.157	−0.338	−0.251	−0.077	−0.024
	(0.484)	(0.416)	(0.186)	(0.204)	(0.743)	(0.920)
2020*Upload speed	−0.013	−0.001	0.018	0.012	−0.010	−0.013
	(0.776)	(0.973)	(0.686)	(0.692)	(0.775)	(0.679)
County-level Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,478	1,286	1,286	1,286	1,286	1,286
R^2	0.001	0.001	0.002	0.002	0.001	0.001

Note: p-values from wild bootstrapping in parentheses. *p<0.1; **p<0.05; ***p<0.01

Table B.2: ACT scores results with average internet download speed (dummy)

	Dependent variables:					
	Composite	Composite (no LA)	English	Math	Reading	Science
2020	0.106	−0.164	−0.239	−.179	−0.137	−0.098
	(0.372)	(0.166)	(0.104)	(0.135)	(0.265)	(0.397)
Low quality	−0.049	0.087	−0.022	0.106	0.128	0.138
	(0.837)	(0.727)	(0.960)	(0.662)	(0.612)	(0.537)
2020*Low quality	−0.047	0.082	0.065	−0.058	0.230	0.058
	(0.911)	(0.854)	(0.911)	(0.894)	(0.601)	(0.881)
COVID-19 rate	−0.027***	−0.017***	−0.021***	−0.018***	−0.017***	−0.016***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Constant	20.128***	20.260***	19.468***	20.005***	20.650***	20.457***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	1,475	1,286	1,286	1,286	1,286	1,286
R^2	0.025	0.013	0.012	0.014	0.011	0.011
Adjusted R^2	0.022	0.009	0.009	0.011	0.008	0.008

Note: p-values from wild bootstrapping in parentheses. *p<0.1; **p<0.05; ***p<0.01

Table B.3: Percentage of qualifying AP scores results with average internet download speed (continuous)

	Dependent variable:
	AP scores
2020	4.913*
	(0.053)
2020*Download speed	0.135
	(0.746)
County-level Fixed Effects	Yes
Observations	1,206
R^2	0.030

Note: p-values from wild bootstrapping in parentheses.
*p<0.1; **p<0.05; ***p<0.01

Table B.4: Percentage of qualifying AP scores results with average internet download speed (dummy)

	Dependent variable:
	AP scores
2020	6.452***
	(0.000)
Low quality	−0.614
	(0.868)
2020*Low quality	2.419
	(0.704)
COVID-19 rate	−0.303***
	(0.001)
Constant	46.489***
	(0.000)
Observations	1,206
R^2	0.038
Adjusted R^2	0.035

Note: p-values from wild bootstrapping in parentheses.
*p<0.1; **p<0.05; ***p<0.01

Table B.5: Graduation rate results
with average internet download speed (continuous)

	Dependent variable:
	Graduation rate
2020	−1.120
	(0.630)
2020*Download speed	0.857**
	(0.050)
County-level Fixed Effects	Yes
Observations	2,543
R^2	0.007

Note: p-values from wild bootstrapping in parentheses.
*p<0.1; **p<0.05; ***p<0.01

Table B.6: Graduation rate results
with average internet download speed (dummy)

	Dependent variable:
	Graduation rate
2020	3.9131***
	(0.000)
Low quality	0.129
	(0.957)
2020*Low quality	−1.581
	(0.687)
COVID-19 rate	0.122**
	(0.010)
Constant	80.951***
	(0.000)
Observations	2,540
R^2	0.008
Adjusted R^2	0.007

Note: p-values from wild bootstrapping in parentheses.
*p<0.1; **p<0.05; ***p<0.01

Table B.7: ACT scores results with median internet upload speed (continuous)

	Dependent variables:					
	Composite	Composite (no LA)	English	Math	Reading	Science
2020	0.152*	−0.170***	−0.252***	−0.201***	−0.153**	−0.078
	(0.052)	(0.007)	(0.000)	(0.002)	(0.039)	(0.190)
2020*Upload speed	−0.016	0.003	0.006	0.007	0.007	−0.007
	(0.457)	(0.872)	(0.779)	(0.696)	(0.727)	(0.714)
County-level Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,478	1,286	1,286	1,286	1,286	1,286
R^2	0.000	0.002	0.002	0.002	0.001	0.001

Note: p-values from wild bootstrapping in parentheses. *p<0.1; **p<0.05; ***p<0.01

Table B.8: ACT scores results with median internet upload speed (dummy)

	Dependent variables:					
	Composite	Composite (no LA)	English	Math	Reading	Science
2020	0.128	−0.122	−0.194	−0.151	−0.070	−0.066
	(0.396)	(0.423)	(0.298)	(0.335)	(0.652)	(0.656)
Low quality	−0.045	−0.025	−0.057	0.016	−0.021	−0.017
	(0.716)	(0.834)	(0.712)	(0.901)	(0.875)	(0.872)
2020*Low quality	−0.059	−0.089	−0.099	−0.074	−0.131	−0.068
	(0.807)	(0.724)	(0.736)	(0.779)	(0.608)	(0.789)
COVID-19 rate	−0.027***	−0.018***	−0.021***	−0.018***	−0.017***	−0.016***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Constant	20.145***	20.277***	19.492***	20.005***	20.669***	20.474***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	1,475	1,286	1,286	1,286	1,286	1,286
R^2	0.025	0.013	0.013	0.014	0.011	0.011
Adjusted R^2	0.023	0.010	0.010	0.011	0.008	0.008

Note: p-values from wild bootstrapping in parentheses. *p<0.1; **p<0.05; ***p<0.01

Table B.9: Percentage of qualifying AP scores results with median internet upload speed (continuous)

	Dependent variable:
	AP scores
2020	6.563***
	(0.000)
2020*Upload speed	−0.328
	(0.157)
County-level Fixed Effects	Yes
Observations	1,206
R^2	0.029

Note: p-values from wild bootstrapping in parentheses. *p<0.1; **p<0.05; ***p<0.01

Table B.10: Percentage of qualifying AP scores results with median internet upload speed (dummy)

	Dependent variable:
	AP scores
2020	6.479***
	(0.000)
Low quality	−0.975
	(0.464)
2020*Low quality	0.086
	(0.964)
COVID-19 rate	-0.311***
	(0.000)
Constant	46.947***
	(0.000)
Observations	1,206
R^2	0.039
Adjusted R^2	0.035

Note: p-values from wild bootstrapping in parentheses. *p<0.1; **p<0.05; ***p<0.01

Table B.11: Graduation rate results
with median internet upload speed (continuous)

	Dependent variable:
	Graduation rate
2020	2.172***
	(0.000)
2020*Download speed	0.568**
	(0.014)
County-level Fixed Effects	Yes
Observations	2,543
R^2	0.006

Note: p-values from wild bootstrapping in parentheses.
*p<0.1; **p<0.05; ***p<0.01

Table B.12: Graduation rate results
with median internet upload speed (dummy)

	Dependent variable:
	Graduation rate
2020	4.953***
	(0.000)
Low quality	4.539***
	(0.000)
2020*Low quality	−2.957*
	(0.082)
COVID-19 rate	0.112**
	(0.019)
Constant	79.277***
	(0.000)
Observations	2,540
R^2	0.015
Adjusted R^2	0.014

Note: p-values from wild bootstrapping in parentheses.
*p<0.1; **p<0.05; ***p<0.01